An Evaluation of Harmonic Isolation Techniques for Three Phase Active Filtering

A Thesis submitted in partial fulfilment of the
requirements for the Degree of Master of Engineering
(Electrical and Electronic)

at the
University of Canterbury
Christchurch, New Zealand

by
David Mark Edward Ingram, B.E. (Hons)

March 1998

Abstract

Recent advances in power electronics have lead to the wide spread adoption of advanced power supplies and energy efficient devices. This has lead to increased levels of harmonic currents in power systems, degrading the performance of electrical machinery and interfering with telecommunication services. Active filters provide a solution to these problems by compensating for the distorted currents drawn by non-linear loads. Optimal methods for controlling these active filters have been determined by computer simulation and experimental implementation.

Methods used for isolating the harmonic content of an unbalanced three phase load current were compared by computer simulations. A technique based on the Fast Fourier Transform (FFT) was developed as part of this work and shown to perform favourably. Notch Filtering, Sinusoidal Subtraction, Instantaneous Reactive Power Theory, Synchronous Reference Frame and Fast Fourier Transform methods were simulated. The methods shown to be suitable for compensation of three phase unbalanced loads were implemented in a Digital Signal Processor to evaluate true performance. These methods were Notch Filtering, Sinusoidal Subtraction, Fast Fourier Transform, and a High Pass Filter based method.

A completely digital hysteresis current controller for a three phase active filter inverter has been developed and implemented with a Field Programmable Gate Array. This controller interfaces directly to a digital signal processor and is resistant to electromagnetic interference.

Results from the experimental hardware verified that the active filter model used for simulation is accurate, and may be used for further development of harmonic isolation methods. A technique using notch filtering gives the best performance for steady loads, with the FFT based technique giving the most flexible operation for a range of load current characteristics. Novel use of the FFT based harmonic isolation technique allows selective cancellation of individual harmonics, with particular application to multiple shunt filters connected in parallel.

Acknowledgments

Firstly I would like to thank my supervisor, Dr Simon Round, for all of his support and guidance throughout this work. I am especially grateful for the financial support for myself and the Department provided by David Dawson of Metalect Industries (NZ) Ltd. Without Mr Dawson's kind generosity and supply of a powerful computer this work could not have been completed. My thanks to Telecom New Zealand for the Telecom Fellowship in Telecommunications Engineering.

The technical assistance provided by Shayne Crimp, Ken Smart and Ron Battersby is greatly appreciated. This work would not have been completed without their down-to-earth wisdom and 'real world' experience. Many thanks to Paul Sinclair for his help in the long-distance printing and submission of the thesis. The weekly lab meetings run by Dr Richard Duke helped clear the air and provided plenty of fresh ideas when the going got tough. Thanks to Paul Sinclair, Adam Taylor and Hamish Laird for your suggestions.

My parents deserve special mention for the support provided to me while I was at university, and for the many years before that. My final thanks go to my wife, Wendy, for all her support throughout my degree. The endless hours spent proofreading are really appreciated.

Table of Contents

List of Figures

List of Tables

Nomenclature

Acronyms

AC	Alternating Current
ADC	Analogue to Digital Converter
CF	Crest Factor
DAC	Digital to Analogue Converter
DC	Direct Current
DSP	Digital Signal Processor
EMI	Electromagnetic Interference
EPLD	Erasable Programmable Logic Device
EPROM	Erasable Programmable Read Only Memory
FFT	Fast Fourier Transform
FIR	Finite Impulse Response
FPGA	Field Programmable Gate Array
IC	Integrated Circuit
IFFT	Inverse Fast Fourier Transform
IGBT	Insulated Gate Bipolar Transistor
IIR	Infinite Impulse Response
ISA	Industry Standard Architecture
PC	Personal Computer
PCB	Printed Circuit Board
PWM	Pulse Width Modulation
RAM	Random Access Memory
RMS	Root Mean Square
SRM	Sample Rate Multiplier
THD	Total Harmonic Distortion
VCO	Voltage Controller Oscillator

Symbols

I_C	Compensating Current
I'_C	Compensating Current Signal
I_L	Load Current
I'_L	Load Current Signal
I_S	Supply Current
V_{DC}	DC voltage

Chapter 1
Introduction

In previous decades electricity has predominantly been used to power electric motors, resistive heating and incandescent lighting. Now, increasing amounts of electricity are being used by electronic loads: consumer electronics such as televisions, industrial electronics such as variable speed motor drives, and high efficiency lighting. These modern devices are becoming widely adopted due to their energy efficient operation. Much has been made of the environmental impact of energy saving devices, but there is another side to the use of modern electronic equipment. While trying to save the Earth's environment, pollution has been added to the electromagnetic and power environments. Modern electronic equipment generally draws non-sinusoidal currents from the AC supply. These harmonic currents are not naturally present, but "pollute" and have a large impact on power systems.

The use of electronic loads is steadily increasing, with about 50%–60% of all electric power in industrialised countries flowing through some type of power electronic system (Redl, Tenti and van Wyk, 1997). By the year 2000 it is expected that 60% of electricity will be passing through non-linear loads (Bernard, 1997). The increasing level of harmonic currents being drawn from the power systems of the world is leading to greater harmonic voltage levels in power distribution systems. An example is Switzerland, where the voltage harmonic content of the 230V distribution system has increased from 3.6% in 1979 to 4.7% in 1991 (Redl *et al.*, 1997).

Problems caused by electromagnetic pollution were first noted in the 1930s when unintentional interference with radio transmissions was noticed. In 1933 CISPR (Comité International Spécial des Perturbations Radioélectriques) was formed to perform regulatory work in Europe. Since then the field of Electromagnetic Compatibility (EMC) has developed along with a large range of standards from CISPR itself, the International Electrotechnical Commission (IEC) and the International Organisation for Standardisation (ISO).

This chapter covers the generation of harmonic currents and their detrimental effects. Several techniques exist for removing harmonic currents from power system and these will be introduced, with particular attention to active filtering. Current standards covering the acceptable limits of harmonic currents are discussed. Special features of this thesis will be presented, along with a summary of the remaining chapters.

1.1 Harmonics Currents and the Power System

1.1.1 Generation of Harmonic Currents

Power electronics devices have created harmonic currents from the beginning of the twentieth century when the use of mercury arc rectifiers began. Advanced control techniques such as pulse width modulation did not occur until power semiconductors were introduced, and although these allowed more flexible operation, these new devices generated more harmonics (van Wyk, 1993).

The era of power semiconductors was started by the thyristor and was followed by the triac, gate turn-off thyristor (GTO), bipolar junction transistor (BJT), power MOSFET, insulated gate bipolar transistor (IGBT) and other more advanced devices (Bose, 1992). These modern semiconductors have higher switching frequencies and have lead to the widespread industrial use of pulse width modulation (PWM) for motor control and power supply regulation (Bose, 1992). PWM is not naturally commutated and so leads to higher levels of harmonic currents than were present with the older mercury arc rectifiers (van Wyk, 1993).

A change in design philosophy is also affecting the harmonic generating nature of modern devices. In the past, equipment was under-rated and over designed. Now, power devices and other pieces of equipment are more critically designed, pushing devices to their operating limits. In some designs magnetic materials are operated in their non-linear regions to reduce manufacturing costs, increasing the distortion of the current drawn by these devices (IEEE Working Group on Power System Harmonics, 1983).

1.1.2 The Effect of Harmonic Currents

Although the effects of harmonic currents and voltages in power systems are varied, the result is that harm is often caused to other equipment connected to the same power system. Power system engineers must consider the effect of harmonic currents on a wide variety of equipment (Rice, 1986). Susceptible equipment may be installed at the site where harmonics are generated, but more often than not, the affected electrical machinery is elsewhere. The following effects are caused by harmonics in the power system (IEEE Working Group on Power System Harmonics, 1983):

- excessive losses and heating of induction and synchronous machines
- mechanical oscillation of induction and synchronous machines
- overvoltages and excessive currents from harmonic resonances

- dielectric breakdown of insulated cables from harmonic overvoltages
- capacitor bank failure from dielectric breakdown
- unstable operation of firing circuits based on zero voltage crossing detection or latching
- interference with ripple control and power line carrier systems
- errors in induction (Ferraris) kilowatt-hour meters
- inductive interference to telecommunication systems

The above list gives the effects that harmonic currents and voltages have directly on electrical and electronic equipment, leading to defects that cost the operators of the equipment and any third parties dependent on that equipment. Harmonic current increases are tied to economic effects: cost control, quality control, efficiency and productivity. The economic disadvantage is the expense of 'cleaning up' the power—this costs at least US$1.2 billion per annum and is increasing (Bernard, 1997). Another detrimental effect is the reduced serviceable life of electrical machinery. The Canadian Electrical Association has calculated reductions in service life due to the presence of harmonics in power systems of 33% for single phase machines, 18% for three phase machines and 5% for transformers and universal machines (Bernard, 1997). Harmonics do cause problems and so there are compelling reasons to remove them from the power system.

1.2 Elimination of Harmonic Currents

Two methods of removing harmonic currents from power systems are commonly used. The first, and oldest, method is the use of passive tuned filters to shunt harmonic currents to ground. A more modern technique, and the method that will be discussed for the rest of this thesis, is shunt active filtering. Methods for removing harmonic voltages are outside the scope of this thesis and will not be discussed.

1.2.1 Shunt Passive Filters

Passive filters are a simple concept, but can be more complicated in reality. The design of passive filters is quite involved and is covered in some detail in texts dealing specifically with harmonics (Arrillaga, Bradley and Bodger, 1985). Two types of passive filters are band pass, for specific harmonics, and high pass, for a range of harmonics. A typical passive filter installation is shown in Figure 1.1. This would have a band pass filter for the fifth and seventh harmonics and a high pass filter for the eleventh harmonic onwards.

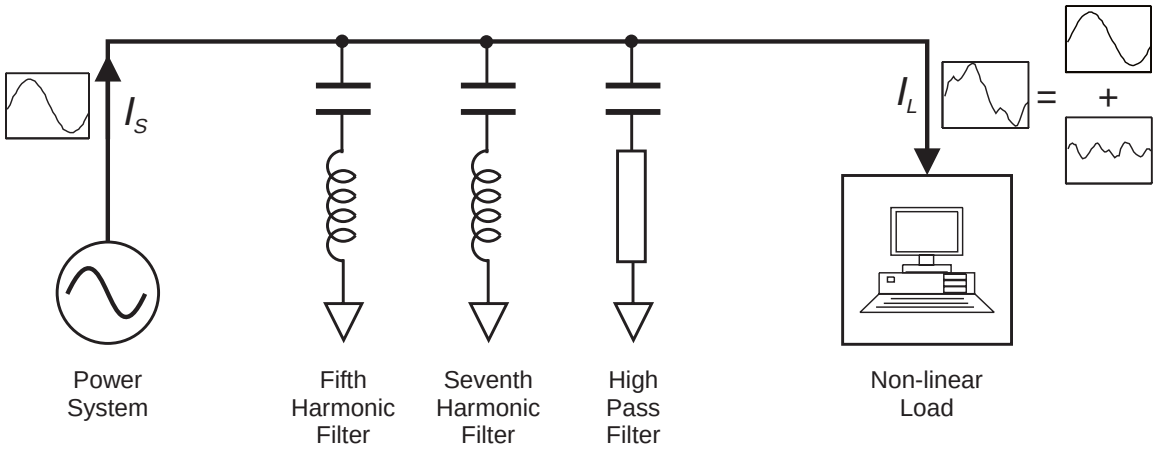

Figure 1.1 Passive filtering of harmonics using band pass and high pass filters.

For the passive filters to be effective, their impedance to ground at harmonic frequencies has to be much lower than that of the power system. This presents problems at low voltage ($400V_{l\text{-}l}$) as the system impedance is very low. Passive filters are typically used on large DC drives, such as dredges and cable cars, which are powered directly from an 11kV supply. Extremely large filters are used at High Voltage DC (HVDC) converter stations to remove the large quantities of harmonic currents generated.

Each band pass passive filter is tuned to a particular frequency, such as 250Hz for the fifth harmonic, and has a finite bandwidth. Should the mains frequency vary, as it does regularly, the effectiveness of the filter will change. This is particularly noticeable with filters for higher order harmonics where the frequency shift Δf is multiplied by the harmonic number. To compensate for the frequency shift a wider bandwidth filter is designed, at the expense of its rejection capability.

Passive filters work by shunting harmonic currents to ground. Unfortunately the filter is not selective about which harmonic currents are shunted to ground—if the harmonic frequency is correct, the filter provides a low impedance path. If other harmonic producing end-users are on the same power feeder, as shown in Figure 1.2, all the harmonics produced will pass through the filter and cause overloading (Arrillaga *et al.*, 1985).

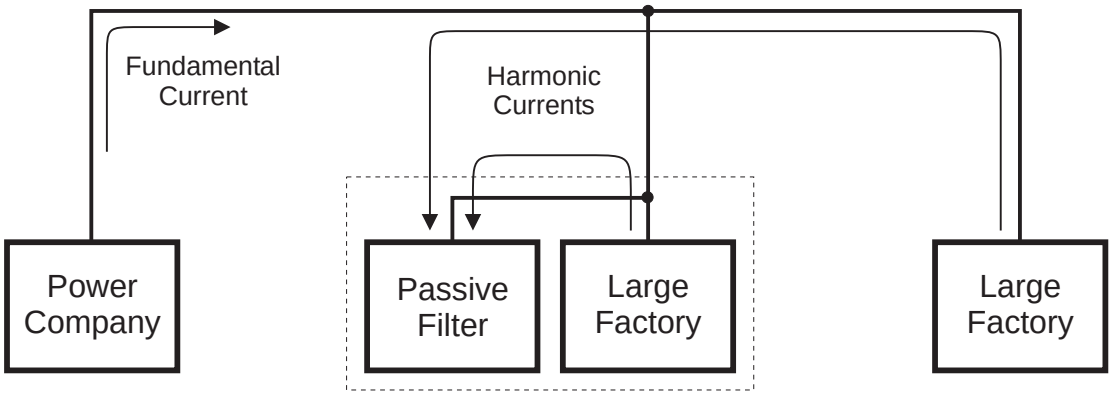

Figure 1.2 Illustration of a passive filter sinking unwanted harmonics to ground.

Shunt passive filters are capacitive below their resonant frequency, so a parallel resonance may occur with inductive loads. Series resonance between the filter capacitance and the supply impedance may cause large levels of harmonic currents to flow from the power system into the shunt filter. Both situations can lead to the destruction of the filter (Fujita and Akagi, 1991). Active filters provide a solution to this and other problems discussed in this section.

1.2.2 Shunt Active Filters

Sasaki and Machida originally proposed the use of shunt active filters as a method of removing current harmonics (Sasaki and Machida, 1971). Recent advances in semiconductor technology have produced high-speed, high-power devices suitable for constructing active filters (Duke and Round, 1993).

Figure 1.3 is a simplified representation of an active filter where the three phases have be drawn as a single line. I_L is the load current drawn by the non-linear load and is sensed by a current transducer. The resulting signal, I'_L, is fed into the digital signal processor (DSP) that performs harmonic isolation. The ‘Harmonic Isolator’ separates the harmonic current signal from the fundamental component of the load current signal so a compensating current can be calculated. A variety of methods for isolating harmonic currents are used and will be discussed in Section 2.2. The compensating current signal, I'_C, is the inverse (a 180° phase shift) of the harmonic current signal. This is used to control a three phase inverter so the required compensating current, I_C, is injected into the power line. This cancels the harmonics drawn from the non-linear load and the resulting supply current, I_S, is ideally sinusoidal.

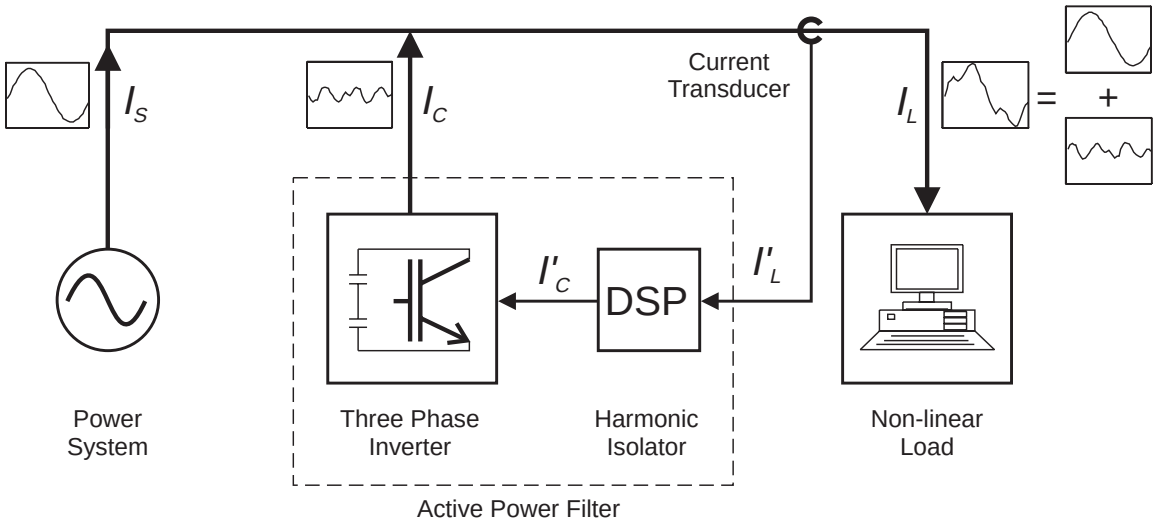

Figure 1.3 Single line diagram of an active filter.

A three phase inverter can generate the harmonic currents without consuming real power other than that required to cover losses in the transistors. The capacitors are used for instantaneous energy storage, and are charged from the power system by adding small levels of fundamental to the compensating current and maintained at a constant DC voltage. A bus voltage controller is used to match the incoming real power with the losses of the power components.

Shunt active filters only remove the harmonic currents associated with the load and therefore do not provide a low impedance to ground for other harmonics in the power system (Duke, Round and Henderson, 1990). As the active filter has no associated passive filters and can therefore be made to track the mains frequency, wide bandwidth filters are avoided (Round, 1992). Figure 1.4 illustrates the ability of an active filter to remove harmonic currents from its associated load while ignoring those from other sites. Figure 1.3 shows the use of a current transducer on the load side of the active filter to determine the load current. The compensating current is calculated from this load current and no other. Any harmonics caused by other sites on the power line remain, but no extra harmonics are added from the site with an active filter. If all sites installed appropriate sized active filters, all harmonics would be eliminated before they could enter the power system.

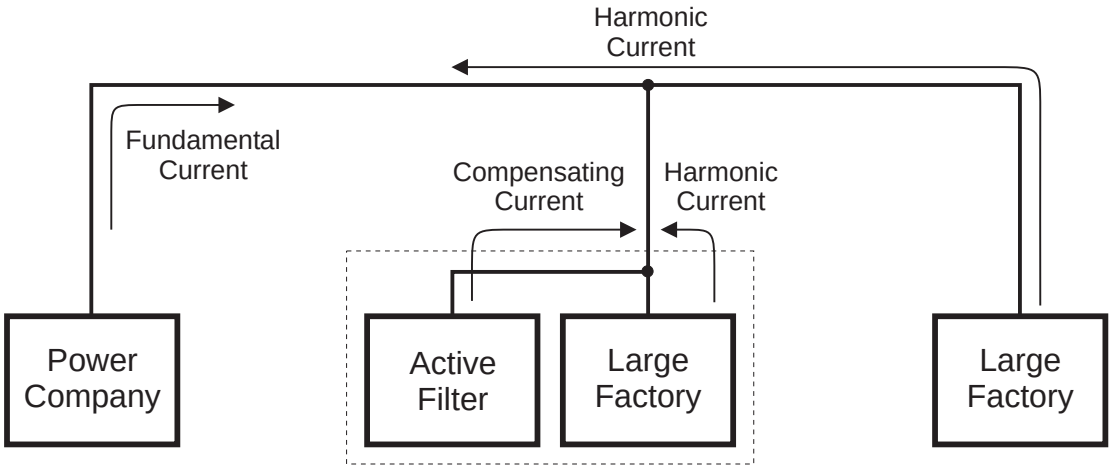

Figure 1.4 Representation of an active filter only compensating the harmonics produced at the site where it is installed.

More than 500 shunt active filters have been used in practical applications since 1981. The market is developing as the price gradually decreases due to the reducing cost of semiconductors and the improvement of power electronics integration (Akagi, 1996).

1.3 Harmonic Standards and Regulations

As the problems caused by harmonics become recognised around the world, standards setting bodies are creating electrical standards that define legal limits for the level of harmonic currents and voltages. The Institute of Electrical and Electronic Engineers has drafted a Recommended Practice (IEEE Std. 519, 1992) that provides limits for harmonic distortion. IEEE Std. 519 limits the current harmonics that can be drawn from the power system. These limits are proportional to the short circuit current ratio and each consumer must limit the current that they draw accordingly (Duffey and Stratford, 1989; Laird, 1992). The aim of the standard is to ensure that voltage harmonic distortion is kept low by limiting the harmonic currents drawn by end users. This ‘standard’ is being rapidly adopted by the electricity utilities in the United States (Bernard, 1997).

The International Electrotechnical Commission (IEC) has a standard, IEC 61000-3-2, that defines harmonic current limits for devices with a current rating less than or equal to 16A. This has been ratified as a Harmonised European Standard, EN 61000-3-2, and as a British Standard (BS EN 61000-3-2, 1995). Unlike its predecessor (IEC 555-2, 1982), no distinction is made between domestic and professional equipment; rack mounted and three phase equipment is specifically mentioned in BS EN 61000-3-2. Electrical equipment is

broken down into several classes, with each class having different harmonic limits. All devices other than portable tools, lighting or equipment having a ‘special wave shape’ is Class A. Class A limits are absolute since maximum harmonic currents are specified. Portable tools form Class B and have one and a half times the maximum permissible harmonic currents of Class A. Class C, lighting equipment, is different in that harmonic current limits are expressed as percentages of the fundamental current. Class D, for equipment with power ratings between 75W and 600W with a special wave shape, expresses the maximum harmonic currents as a percentage of the power rating of the device. The ‘special wave shape’ for Class D has a high peak and is shown in Figure 1.5. The centre line, “M”, coincides with the peak value of the input current. A typical current that would fit this waveform would be from single phase bridge rectifiers with capacitive smoothing. These circuits are often used as the DC source in a switching power supply, such as those found in personal computers (Round and Ingram, 1997).

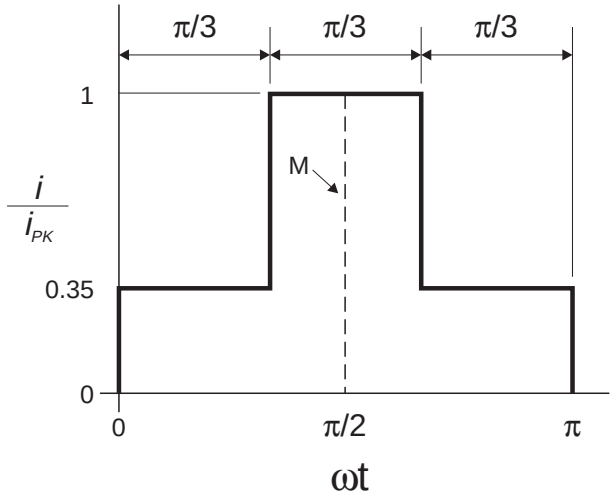

Figure 1.5 Envelope of the input current to define the “special wave shape” that classifies equipment as Class D, from BS EN 61000-3-2.

Harmonics in New Zealand are limited by the Electrical Code of Practise ECP36, but this code of practise only deals with voltage harmonics. Two Australian Standards (AS 2279.1, 1991; AS 2279.2, 1991) deal with harmonic currents in ‘mains supply networks’. AS 2279.2 deals with equipment with power ratings greater than 4.8 kVA, intended exclusively for industrial, professional or commercial purposes, or equipment requiring an electricity supply authority’s consideration before connection. AS 2279.1 deals with the remaining equipment, both single phase 240V two wire and three phase 240/415V three and

four wire apparatus. The current limits in AS 2279.1 are absolute, with the limits being slightly higher than those of Class A in BS EN 61000-3-2.

Individual pieces of equipment can be designed to meet absolute harmonic limits. Harmonic problems arise when many small loads are installed at one site, such as personal computers in a university computer laboratory or large office building. Some progress has been made by specifying harmonic current limits for lighting (Class C) as a percentage of fundamental current (BS EN 61000-3-2, 1995). Until all standards are changed to specify harmonic limits this way, designers have no reason to increase the cost of their products by drawing sinusoidal currents from the power system.

1.4 Overview

The aim of this thesis is to present the results of an investigation into the performance of several harmonic isolation techniques. This work continues on from the digitally controlled single phase active filter developed by Dr S.D. Round (Round, 1992; Round and Duke, 1994). Other active filtering research conducted in the Department of Electrical and Electronic Engineering helped form the initial knowledge base (Henderson, 1989; Laird, 1992).

1.4.1 Novel Work

Several aspects of the active filtering research presented in this thesis are novel and worth noting.

- **Experimental Evaluation of Harmonic Isolation Techniques**

  Other work in this field (Grady, Samotyj and Noyola, 1990; Akagi, 1992; Jou, 1995; Akagi, 1996; Horn, Pittorino and Enslin, 1996) investigated the theoretical performance of different techniques for determining the harmonic content of a load current. This work is the first comparison of harmonic isolation techniques, through both simulation and experimental work, for unbalanced three phase applications.

- **Digitally Controlled Three Phase Active Filter**

  Three phase active filters previously designed and built at the University of Canterbury have used analogue controllers (Round, 1992). This is the first implementation of a digital three phase inverter controller.

- **Completely Digital Control of an Inverter**

  As far as the author is aware, the use of a programmable logic device to implement a hysteresis current controller is a first. Other researchers have used a combination of digital signal processors and programmable logic to control inverters but none have run at the high sampling rates achieved by the controller presented in this thesis.

1.4.2 Scope of Thesis

Definitions of useful measures of power quality are presented in Chapter 2, along with a detailed explanation of the harmonic isolation methods that will be examined in this thesis. The initial evaluation of each method was through simulation, with the method and results presented in Chapter 3.

Electronic hardware used to test the different harmonic isolation methods is described in Chapter 4. A digital controller for a power inverter was developed as part of this work and is described in detail in Chapter 5. A digitally interfaced inverter controller is much more resistant to electromagnetic interference (EMI) than an analogue controller (Ingram and Round, 1997). This makes this controller more suitable for use in a switching inverter where EMI levels can be high.

The test software is described in Chapter 6 and the design and operation of digital filters used in this work presented in Chapter 7. Chapter 8 details the experimental results and a comparison with the simulation results of Chapter 3 is given in Chapter 9, along with future work. The author’s conclusions are presented in Chapter 10.

Chapter 2
Background

2.1 Power Quality Definitions

Several expressions are used in this thesis to describe the harmonic content of a current waveform. The first is the Root Mean Square (RMS) which gives an indication of the magnitude of a voltage or current and takes into account the fundamental and all the harmonics. An expression for the RMS value of a current based on its frequency domain components is given in Eqn. (2.1), where I_h is the h^{th} harmonic. When a current waveform is sampled in time, these samples can be used to calculate the RMS value of the current. Eqn. (2.2) gives the expression for the calculation of the RMS based upon N samples in the time domain.

$$I_{RMS} = \sqrt{\sum_{h=1}^{\infty} I_h^2} \tag{2.1}$$

$$I_{RMS} = \frac{\sqrt{\sum_{n=1}^{N} I_n^2}}{N} \tag{2.2}$$

Total Harmonic Distortion (THD) is the most commonly used index of the distortion present in a waveform (Stanislawski *et al.*, 1997). One of the more commonly used definitions of THD, and the one used in New Zealand, is given by Eqn. (2.3).

$$THD = \frac{\sqrt{\sum_{h=2}^{\infty} I_h^2}}{I_1} \tag{2.3}$$

The peak current of a distorted waveform will not necessarily be $\sqrt{2}$ times the RMS value of the current, as this only holds true if the current is purely sinusoidal. The Crest Factor (CF) is the ratio of the peak current to the RMS current, as shown by Eqn. (2.4). A highly distorted current may have a CF as high as five, especially the currents drawn by single phase switching power supplies.

$$CF = \frac{I_{Peak}}{I_{RMS}} \tag{2.4}$$

2.2 Harmonic Isolation Methods

Harmonic isolation methods are grouped into three classifications; load current, supply current and supply voltage. These each have their own characteristics, particularly their response to transients. Load current and supply current methods are most suitable for shunt active filters installed near a harmonic producing load such as a commercial office building. Voltage detection is the most suitable technique to be used for an active filter that is going to be connected in series as part of a Unified Power Flow Controller (Akagi, 1996) and will not be discussed further here.

The main difference between active filters that use load current sensing or supply current sensing is the location of the current transducer. Harmonic current signals from the transducer signal are analysed and used to control the active filter. Figure 2.1 shows the location of current transducers for these two types of active filtering. Methods that use the load current in the calculation of the compensating current for a shunt active filter examine the current drawn by the non-linear load using a current transducer ‘downstream’ from the active filter. Several methods exist for isolating the harmonics from a load current and will be discussed in more detail in the following sections. These techniques are Notch Filtering, Sinusoidal Subtraction, Instantaneous Reactive Power Theory, Synchronous Reference Frame and the Fast Fourier Transform. The supply current can also be used in the calculation of the compensating current for an active filter. When a supply current method is used, the current transducer is placed ‘upstream’ of the active filter. The objective of active filtering is to remove harmonic current from the supply current, leaving a pure sinusoid at the fundamental frequency. One technique is to control the active filter’s inverter so that the supply current, I_S, is sinusoidal.

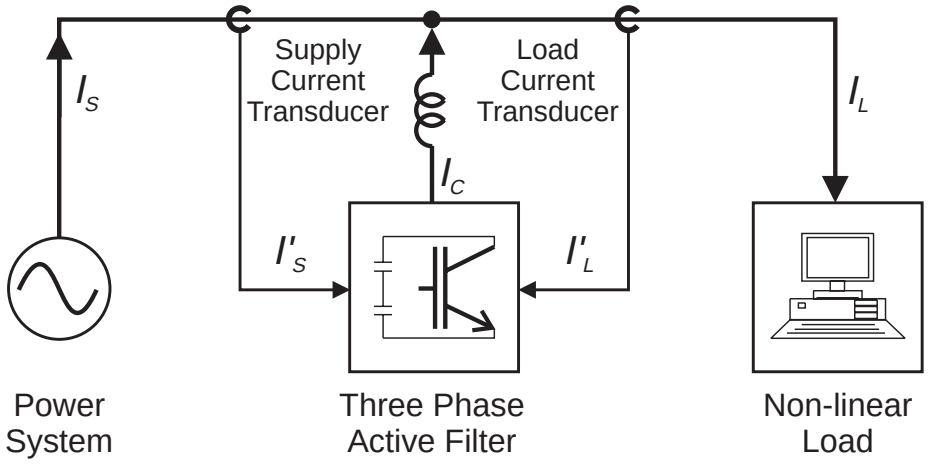

Figure 2.1 Connections for supply side and load side active filtering.

Results presented in this thesis relate to load current based harmonic isolation methods. The majority of publications to date on active filter control use load current techniques (Grady *et al.*, 1990) and provide an established knowledge base for the work presented here. A load current active filter controller takes the load current signal, I'_L (from Figure 2.1), and uses it to compute the compensating current signal, I'_C. The power amplifier, which is discussed in more detail in §4.5, takes I'_C, and injects the true compensating current, I_C, into the power system. This is shown in a more diagrammatic form in Figure 2.2. Each method discussed in the following sections takes I'_L as an input and generates I'_C as an output.

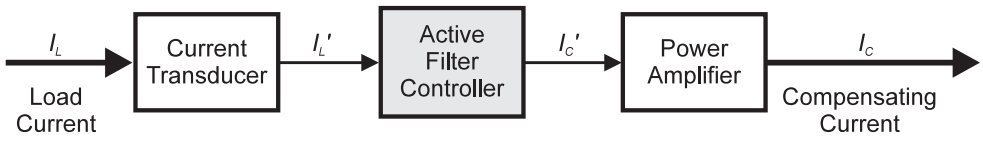

Figure 2.2 Schematic of an active filter with a generic controller.

2.2.1 Notch Filtering

In this method the load current is filtered by a notch filter, which removes the fundamental while retaining the harmonic components. An active filter that uses a notch filter on each of the three phases to isolate the harmonic current from the load current has the

advantage of being able to cope with unbalanced three-phase loads (Quinn, Mohan and Mehta, 1993).

A notch filter, or ‘band stop filter’, removes from a signal a band of frequencies. The complement of this filter is the ‘band pass filter’. Some publications have reported use of true band stop filters to isolate harmonic currents (Quinn *et al.*, 1993), whereas others utilise a band pass filter in combination with a subtracter (Choe and Park, 1988), as shown in Figure 2.3.

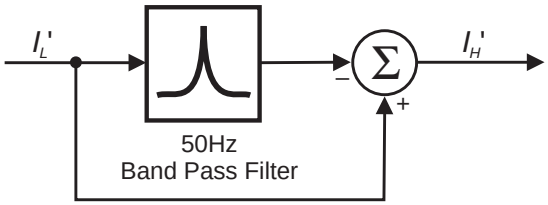

Figure 2.3 Alternative implementation of a band stop filter.

Figure 2.4 shows the block diagram for an active filter that uses a notch filter. The load current is filtered to leave the harmonic currents. These are then effectively subtracted from the load current by injecting into the power line with a 180° phase shift.

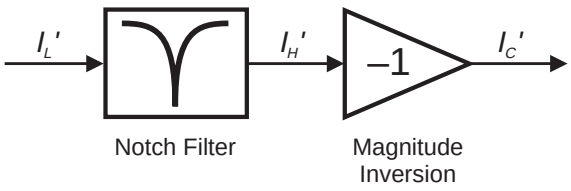

Figure 2.4 Block diagram for a notch filter based active filter.

2.2.2 Sinusoidal Subtraction

This method artificially synthesises a sinusoid of the same magnitude and phase as the load current fundamental and was originally developed by K. Henderson (Henderson, 1989). This synthetic sinusoid is subtracted from the load current, leaving the harmonics. Figure 2.5 shows the block diagram for the sinusoidal subtraction method.

The load current is low pass filtered to yield the magnitude of the 50Hz component that is peak detected every half cycle, but with a phase shift. This gives the magnitude of a

sinewave to be subtracted during the next half cycle. A band pass filter extracts the phase information from the non-linear load current. This has a slow transient response and is not suitable for changing the magnitude of the synthetic sinewave. Having a fast transient response, good harmonic rejection and no phase shift is not possible for a single filter. A combination of a low pass and a band pass filter achieves these requirements. A phase locked loop provides the frequency signal for the sinewave generator to produce a sinewave from a ROM look-up table.

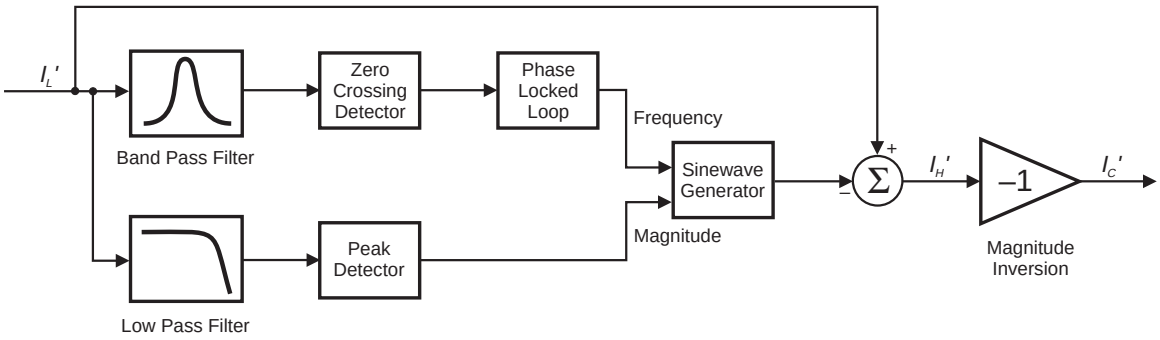

Figure 2.5 Block diagram of the Sinusoidal Subtraction harmonic isolation method.

2.2.3 Instantaneous Reactive Power Theory

Instantaneous Reactive Power Theory (IRPT) uses the Park Transform, given in Eqn. (2.5), to generate two orthogonal rotating vectors (α and β) from the three phase vectors (a, b and c). This transform is applied to the voltage and current, with the symbol *x* used to represent *v* or *i*. It should be noted that IRPT assumes that the three phase load is balanced.

$$\begin{bmatrix} x_\alpha \\ x_\beta \end{bmatrix} = \sqrt{\frac{2}{3}} \begin{bmatrix} 1 & -\frac{1}{2} & -\frac{1}{2} \\ 0 & \frac{\sqrt{3}}{2} & -\frac{\sqrt{3}}{2} \end{bmatrix} \begin{bmatrix} x_a \\ x_b \\ x_c \end{bmatrix} \tag{2.5}$$

The supply voltage and load current are transformed into α-β quantities. The instantaneous active and reactive powers, *p* and *q*, are calculated from the transformed voltage and current as given in Eqn. (2.6). The two powers have a DC component, $\bar{p}$ and $\bar{q}$, and an AC component, $\tilde{p}$ and $\tilde{q}$.

$$\begin{bmatrix} p = \bar{p} + \tilde{p} \\ q = \bar{q} + \tilde{q} \end{bmatrix} = \begin{bmatrix} v_\alpha & v_\beta \\ -v_\beta & v_\alpha \end{bmatrix} \begin{bmatrix} i_\alpha \\ i_\beta \end{bmatrix} \tag{2.6}$$

By looking at instantaneous powers, the harmonic content can be visualised as a ripple upon a DC offset representing the fundamental power. By removing the DC offset with a suitable high pass filter (Pahmer, Capolino and Henao, 1994) and then performing the Inverse Park Transform the harmonic current can be determined (Akagi, Nabae and Atoh, 1986). Figure 2.6 shows the block diagram for an active filter based on Instantaneous Reactive Power Theory. Filtering the instantaneous active and reactive powers leaves the AC components. The compensating currents are calculated by taking the inverse of Eqn. (2.6), as shown by Eqn. (2.7).

$$\begin{bmatrix} i'_\alpha \\ i'_\beta \end{bmatrix} = \frac{1}{v_\alpha^2 + v_\beta^2} \begin{bmatrix} v_\alpha & -v_\beta \\ v_\beta & v_\alpha \end{bmatrix} \begin{bmatrix} \tilde{p} \\ \tilde{q} \end{bmatrix} \tag{2.7}$$

The inverse Park Transform is applied to i'_α and i'_β and this gives the harmonic currents in standard three-phase form, shown in Eqn. (2.8).

$$\begin{bmatrix} i_a \\ i_b \\ i_c \end{bmatrix} = \sqrt{\frac{2}{3}} \begin{bmatrix} 1 & 0 \\ -\frac{1}{2} & \frac{\sqrt{3}}{2} \\ -\frac{1}{2} & -\frac{\sqrt{3}}{2} \end{bmatrix} \begin{bmatrix} i'_\alpha \\ i'_\beta \end{bmatrix} \tag{2.8}$$

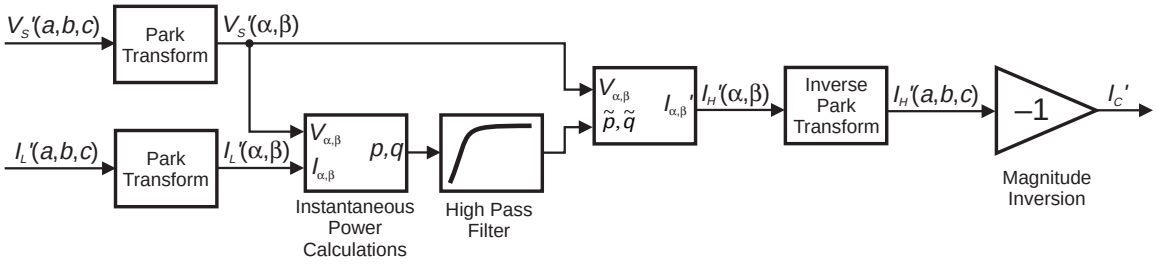

Figure 2.6 Block diagram for an IRPT based controller.

If four-wire three phase systems are to be considered, a more complex control strategy is required. Modifications to the Instantaneous Reactive Power Theory have been proposed to take into account zero sequence terms (Nabae and Tanaka, 1996; Aredes, Häfner and

Heumann, 1997). Eqn. (2.6) has been rewritten to include a zero sequence term, and is given by Eqn. (2.9)

$$\begin{bmatrix} p_0 \\ p \\ q \end{bmatrix} = \begin{bmatrix} v_0 & 0 & 0 \\ 0 & v_\alpha & v_\beta \\ 0 & -v_\beta & v_\alpha \end{bmatrix} \begin{bmatrix} i_0 \\ i_\alpha \\ i_\beta \end{bmatrix} \tag{2.9}$$

Two of these proposals are “Sinusoidal Source Current” and “Constant Source Instantaneous Power”. Both result in artificial balancing of the three phase supply currents because zero sequence power cannot be decomposed into ac and dc termns (Aredes and Watanabe, 1995). Load balancing is undesirable because large amounts of energy are transferred between phases, increasing operating losses in the active filter’s power inverter.

2.2.4 Synchronous Reference Frame

Bhattacharya *et al.*(Bhattacharya, Divan and Banerjee, 1991) proposed using the DQ transform, given in Eqn. (2.10), which changes the three conventional rotating phase vectors into direct (D), quadrature (Q) and zero (0) vectors. The fundamental component for each is now a dc value with harmonics appearing as ripple.

$$\begin{bmatrix} i_0 \\ i_d \\ i_q \end{bmatrix} = \frac{2}{3} \begin{bmatrix} \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ \cos(\omega t) & \cos\left(\omega t - \frac{2}{3}\pi\right) & \cos\left(\omega t + \frac{2}{3}\pi\right) \\ -\sin(\omega t) & -\sin\left(\omega t - \frac{2}{3}\pi\right) & -\sin\left(\omega t + \frac{2}{3}\pi\right) \end{bmatrix} \begin{bmatrix} i_a \\ i_b \\ i_c \end{bmatrix} \tag{2.10}$$

Harmonic isolation of the DQ transformed current signal is achieved by removing the DC offset with a high pass filter. Figure 2.7 illustrates the block diagram of the DQ active filter. Voltage information is not required for a Synchronous Reference Frame (SRF) based controller.

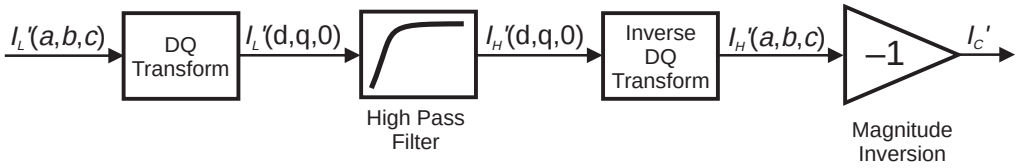

Figure 2.7 Block diagram of the SRF based active filter controller.

As with the IRPT method, SRF based filtering cannot cope with load imbalances. When the load is unbalanced there will be a 100Hz ripple on the D, Q and 0 terms. A 100Hz ripple is also present if third harmonic currents are present (the 150Hz is translated down to 100Hz). The source of the ripple cannot be determined if the load current contains triplen harmonics and is unbalanced.

2.2.5 Fast Fourier Transform

The Fast Fourier Transform (FFT) takes the sampled load current for one period and calculates the magnitude and phase of the frequency components. Figure 2.8(a) shows the 40 time domain samples used as the input to the FFT and Figure 2.8(b) shows the frequency spectrum generated from the FFT output. The spacing between points in the output spectrum is the inverse of the length of the time sample used. In this example, the length of the sample is 20ms and therefore the frequency resolution, Δf, will be 50Hz.

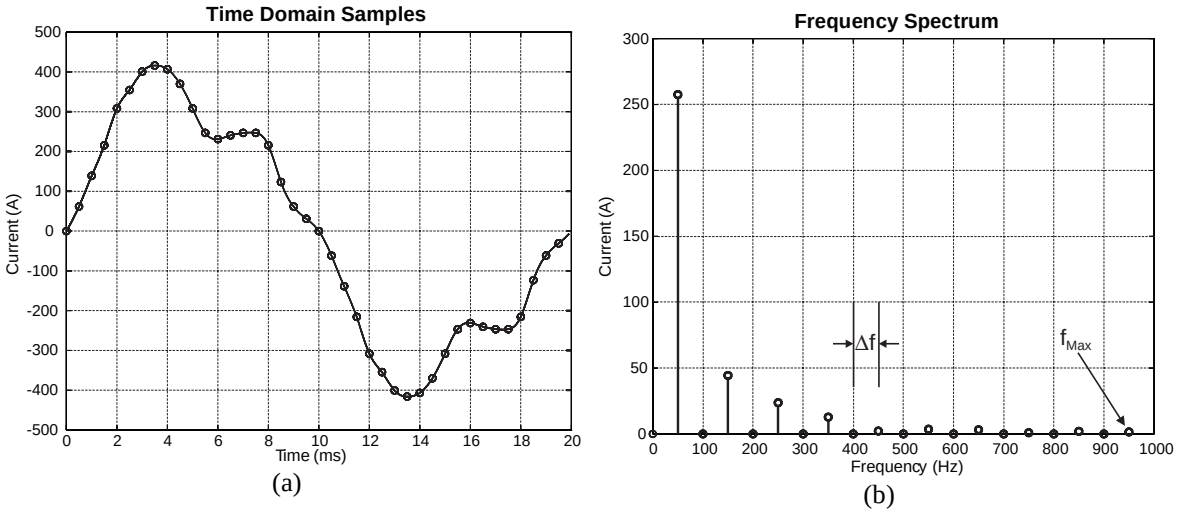

Figure 2.8 The time domain samples (a) of the input waveform used to generate the spectrum in (b).

Each element in the frequency plot is a harmonic since the spacing is 50Hz. The number of harmonics that can be resolved is given by half the number of samples used, minus one. Therefore the higher the number of samples in each cycle of current, the higher the value of f_{Max}. Since forty samples were used, the maximum harmonic that can be resolved in this case is the 19^{th}, or 950Hz. Frequency resolution is increased by sampling for a longer time.

Having more samples during a given interval increases the maximum frequency (and therefore harmonic) that can be resolved.

Removal of the fundamental from the input current is performed by setting the frequency component for 50Hz to zero and then performing the Inverse Fast Fourier Transform (IFFT) on each cycle of current. The IFFT recreates a time domain signal based on the magnitude and phase information of each harmonic (Henderson, 1989). The FFT must be calculated over a complete mains cycle to prevent spectral leakage distorting the output (Arrillaga *et al.*, 1985) which would lead to an incorrect compensating current signal being generated.

A frequency domain based harmonic isolation method has advantages over time domain based techniques. The magnitude of the load harmonics is known from the FFT and this allows selective harmonic cancellation to be performed. Manipulating the harmonic magnitudes makes it possible to prevent the cancellation of certain harmonics or to reduce the compensation of individual harmonics.

2.3 Unbalanced Three Phase Systems

In the Introduction it was explained that the active filter was being developed for three phase, four wire applications where the load currents were not balanced. Simulations of harmonic isolation techniques were used to measure the performance of each with unbalanced three phase load currents.

The five methods discussed so far can be described as either three phase or single phase methods. The three phase methods, IRPT and SRF, take into account the three phase relationship of the load current. Single phase methods (Notch Filtering, Sinusoidal Subtraction and the Fast Fourier Transform) use a single current, requiring three independent harmonic isolators be connected in parallel for a three phase active filter.

2.3.1 Three Phase Methods

The IRPT and SRF techniques each assume that the three phases are balanced in the transformation of three phase currents to a two phase orthogonal representation. When the three phases are balanced, the third current or voltage can be determined from the other two and is an assumption of the three to two variable transformations. If this is not the case, there are three degrees of freedom in the variables, whereas a balanced system has two degrees of freedom. Three degrees of freedom cannot be reduced to two without the loss of information

and therefore when IRPT is performed using unbalanced load currents the active filter will not compensate properly. The SRF method also assumes balanced phases and so the supply current will be incorrect, but not as distorted as with IRPT, in an unbalanced system.

2.3.2 Single Phase Methods in Parallel

Three single phase methods are investigated in this work: Notch Filtering, Sinusoidal Subtraction and the Fast Fourier Transform. Each of these techniques is suitable for use in a single phase active filter. In a three phase active filter, three controllers are used, one for each phase. Implementing the active filter controller with three completely independent harmonic isolation units means no assumptions are made about the three phases. As a result, this type of active filter controller can deal with all types of three phase currents, balanced and unbalanced.

2.4 Summary

Five load current based harmonic isolation methods have been presented in this chapter. Three of these (Notch Filtering, Sinusoidal Subtractions and the Fast Fourier Transform) take a single phase current as their input and one controller is used for each of the three phases. The two remaining methods (Instantaneous Reactive Power Theory and Synchronous Reference Frame) take account of all three phases. IRPT uses load current and system voltage measurements while SRF only uses the three phase load current measurements. Simulations results presented in the next chapter ascertain the suitability of each technique in unbalanced three phase situations.

Chapter 3
Simulation of Harmonic Isolation Techniques

3.1 Simulation with MATLAB and Simulink

Each of the five harmonic isolation methods was simulated in a computer model of an active filter. An identical input current waveform was fed into each active filter simulation and the quality of compensation and transient response determined. Simulink (Mathworks, 1996b) was used to model and simulate the harmonic isolation methods. Simulink is an extension to the matrix manipulation program MATLAB, with a graphical user interface that allows the user to 'draw' the system to be modelled into the computer. Outputs from the model can be graphed as the simulation takes place and can be stored for further processing by MATLAB.

A Simulink model looks very much like a schematic diagram or flow chart. Figure 3.1 is the Simulink block diagram used for simulating the Synchronous Reference Frame active filter controller. The left most block, "Factory", is the source of the non-linear current. The thick lines from "Factory" to "abc2dq0" and "Subtracter" represent vectors and makes connecting multiple signals between blocks much simpler than connecting each signal individually. In this example the vector is a three phase current. The three blocks, "abc2dq0", "AB HPF" (high pass filter) and "dq02abc", are all Simulink models with more detail inside them. Each block has configuration settings accessed by double clicking on the block's icon. The configuration dialogue box for "AB HPF" is shown in the upper left where the user can select the cut-off frequency and the order of the three phase high pass filter. The "Demux" block separates the three phase vector into three signal phase currents to be graphed by "Supply Current Graph" and stored as MATLAB vectors by the "Harmonic Current Output" and "Supply Current Output" blocks. The "Clock" block generates simulation time points to be stored in an MATLAB vector. These time points are then used when plotting current waveforms.

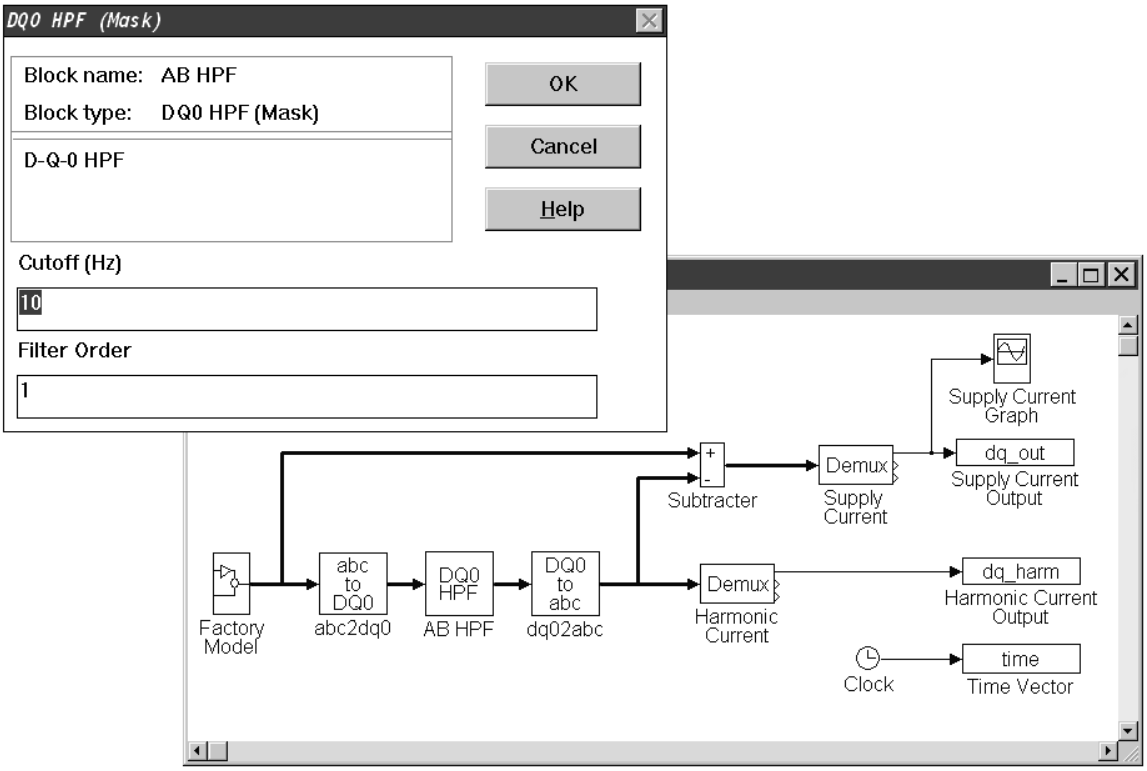

Figure 3.1 The source system for a Simulink simulation.

The Sinusoidal Subtraction and Fast Fourier Transform methods were more easily represented by an algorithm than a block diagram, and were simulated by writing programs in the native language of MATLAB. This is very similar to the BASIC or Pascal languages, yet exploits the powerful matrix and vector operations that MATLAB provides.

3.2 Evaluation of Techniques by Simulation

Thorough testing of each of the harmonic isolation methods was performed as part of this work. The intention was to identify any problems or difficulties with each method before implementation in a digital signal processor (DSP). The work presented in the rest of this chapter shows this was worthwhile as MATLAB simulations acted as a reference for the DSP implementation of each method.

The transient response of each technique was determined by examining at the time taken for each controller to stabilise after a step change in the load and the total harmonic distortion of the resulting supply current. The aim of this work was to implement an active filter that could deal with unbalanced loads, as would be encountered in real situations. Each of the harmonic isolation methods was tested with unbalanced loads and loads that drew

triplen harmonics (harmonic orders that are a multiple of three). Triplen harmonics are unique because they add constructively rather than destructively in three phase systems. This gives a net harmonic current in the neutral wire which invalidates some three phase assumptions, particularly that of Eqn. (3.1), that the IRPT and SRF methods depend upon.

$$I_A + I_B + I_C = 0 \tag{3.1}$$

3.3 Transient Response Performance

A non-linear load current that initially had the characteristic of a fluorescent lighting load was artificially generated in Simulink and is shown in Figure 3.2(a). At a time of 500ms into the simulation the three phase load current from a variable speed motor drive was added to the fluorescent lighting current. This gave a balanced step change in magnitude and harmonic content. A changing harmonic content was intended to represent the switching on of different electrical equipment, rather than changing the current demand of existing machinery.

Current waveforms of the artificial load, I_L, and the resulting supply current, I_S, for each harmonic isolation method are shown in Figure 3.2. The ‘Non-Linear Load’ plot in Figure 3.2(a) appears highly distorted. This load current changes noticeably at 500ms from having predominantly third harmonic currents to one with significant levels of the fifth harmonic. All of the methods, except IRPT, show the same steady-state nearly sinusoidal ‘supply currents’ but with different transient current distortions. The Notch Filtering technique shows very little distortion, but takes six cycles for the current to reach a steady state level. SRF has no supply current distortion and responds to the change quicker than the Notch Filtering method. The Fast Fourier Transform and Sinusoidal Subtraction techniques both experience distortion when the transient occurs, yet reach steady state in two and three cycles respectively. IRPT compensates for the harmonics poorly and does not give a sinusoidal supply current at all. The Total Harmonic Distortion (THD) of each waveform gives a useful measure of the harmonic distortion remaining after compensation.

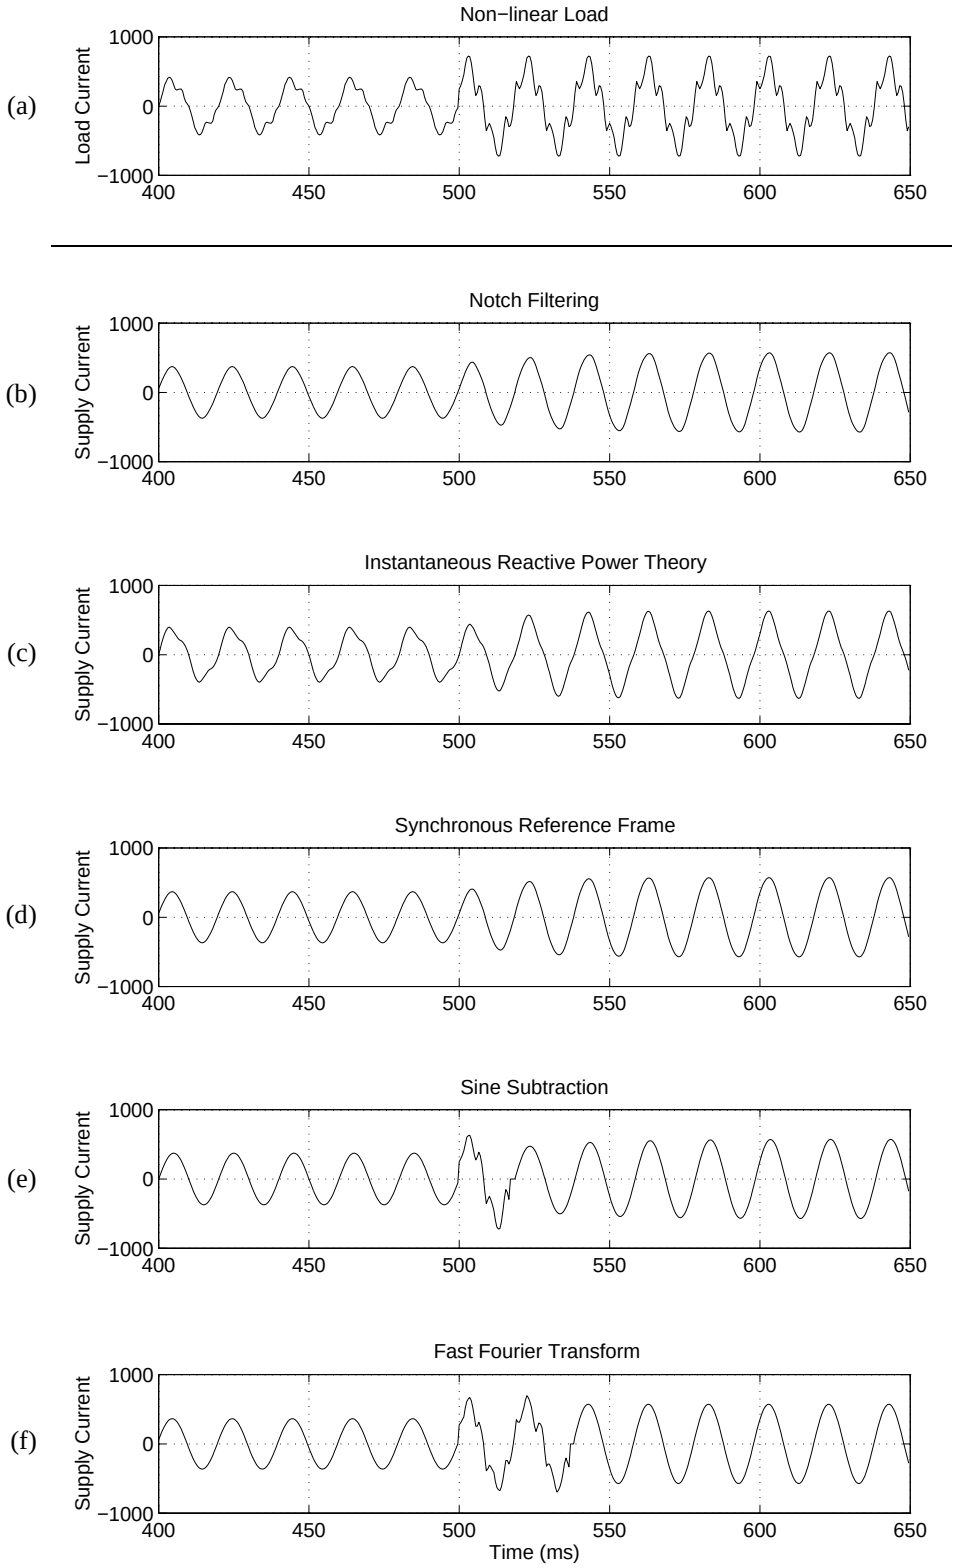

Figure 3.2 Transient response in the time domain from the simulation of each harmonic isolation method.

Total harmonic distortion calculations were made on a sliding window of the previous 20ms of data (for a 50Hz fundamental). This ‘sliding THD’ measured how well each harmonic isolation method dealt with the step changes in current and the change in harmonic content. The sliding THD results for each method are presented in Figure 3.3. The change in harmonic content of the non-linear load at 500ms increased the THD of the load current from 20.3% to 29.6%. Notch Filtering and Synchronous Reference Frame maintained the THD of the supply current at the same level before the transient and after. The slight increases during the transient were due to the non periodic nature of each waveform causing spectral leakage in the Fourier Transform based THD measurement routine. Both the Sinusoidal Subtraction and FFT based harmonic isolation methods gave complete harmonic compensation in the steady state, but did not perform as well as the other techniques during the transient. IRPT compensated better after the transient, once the relative level of the third harmonic was reduced. Table 3.1 lists the THD of the load current and supply current for each harmonic isolation method. THD calculations were made before, during and after the change in load current. The maximum THD recorded during the transient is used to represent current distortion due to step changes.

Table 3.1 Total harmonic distortion of the load and supply currents before, during and after the transient.

	THD (%)		
	Before	**Maximum**	**After**
Load Current	20.3	43.6	29.6
Notch Filtering	2.0	13.8	2.1
IRPT	17.3	22.0	11.0
SRF	1.2	15.8	1.2
Sinusoidal Subtraction	0.0	41.2	0.0
FFT	0.0	38.4	0.0

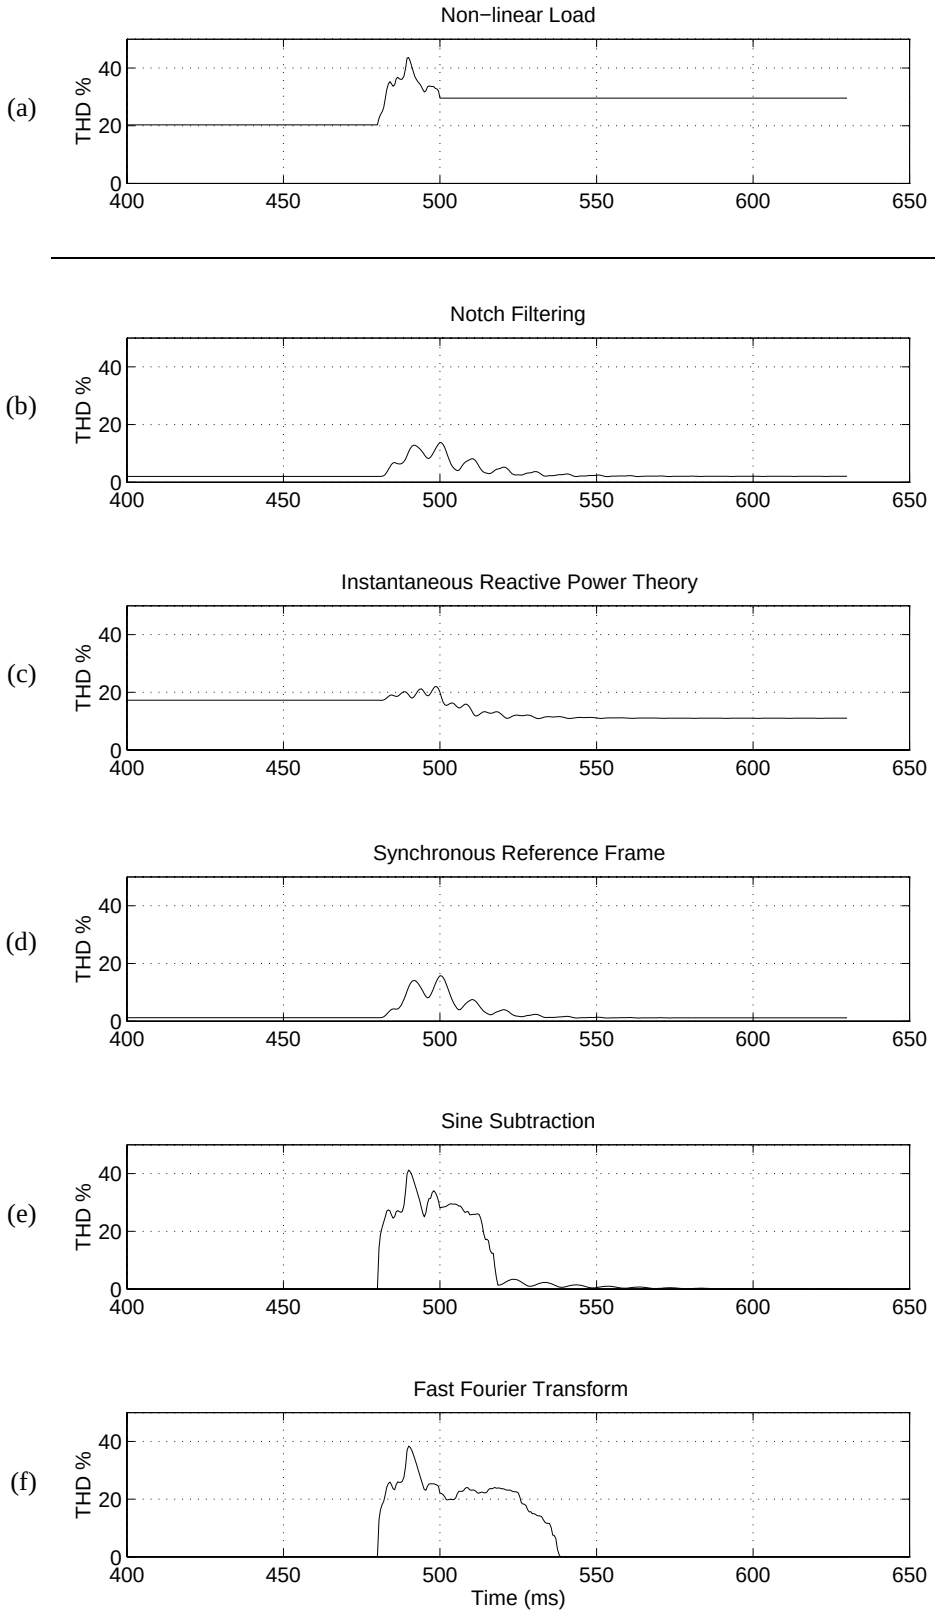

Figure 3.3 Sliding THD of one cycle as each harmonic isolation method responds to a step change.

3.4 Operation With Unbalanced Loads

A Fluke 41 Power Harmonics Analyser was used to sample the current drawn by a personal computer laboratory. A typical load current waveform is illustrated in Figure 3.4(a) with the harmonic spectrum in Figure 3.4(b). It can be seen from Figure 3.4(b) that significant levels of harmonic currents, particularly the third and fifth harmonics, are present. This current was used to generate the load current for simulation in MATLAB. The load currents were artificially unbalanced with Phase A = 100%, Phase B = 50% and Phase C = 150%. In Figure 3.5–Figure 3.7 Phase A is represented by a solid line, Phase B is dashed and Phase C is dotted.

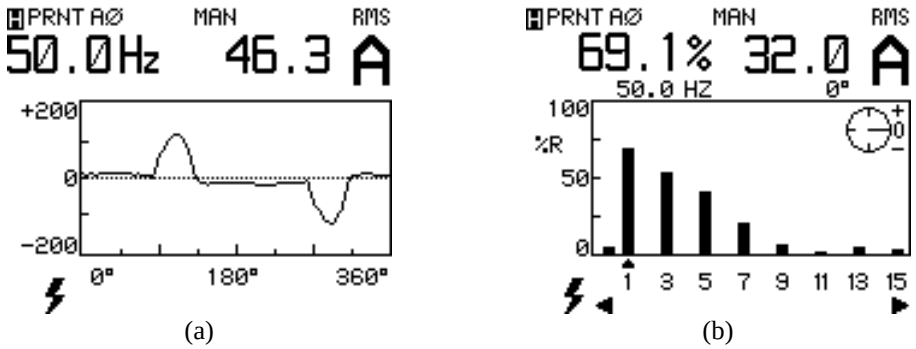

(a) (b)

Figure 3.4 Load current from personal computers used for the unbalanced operation simulation.

The three single phase harmonic determination methods, notch filtering, sinusoidal subtraction and FFT, each performed correctly as these do not make any assumptions about the three-phase currents. The IRPT and SRF methods failed to compensate for the harmonics correctly. Figure 3.5 shows the load current drawn (I_L in Figure 1.3) and the resulting supply currents (I_S in Figure 1.3) for the IRPT, SRF and FFT methods.

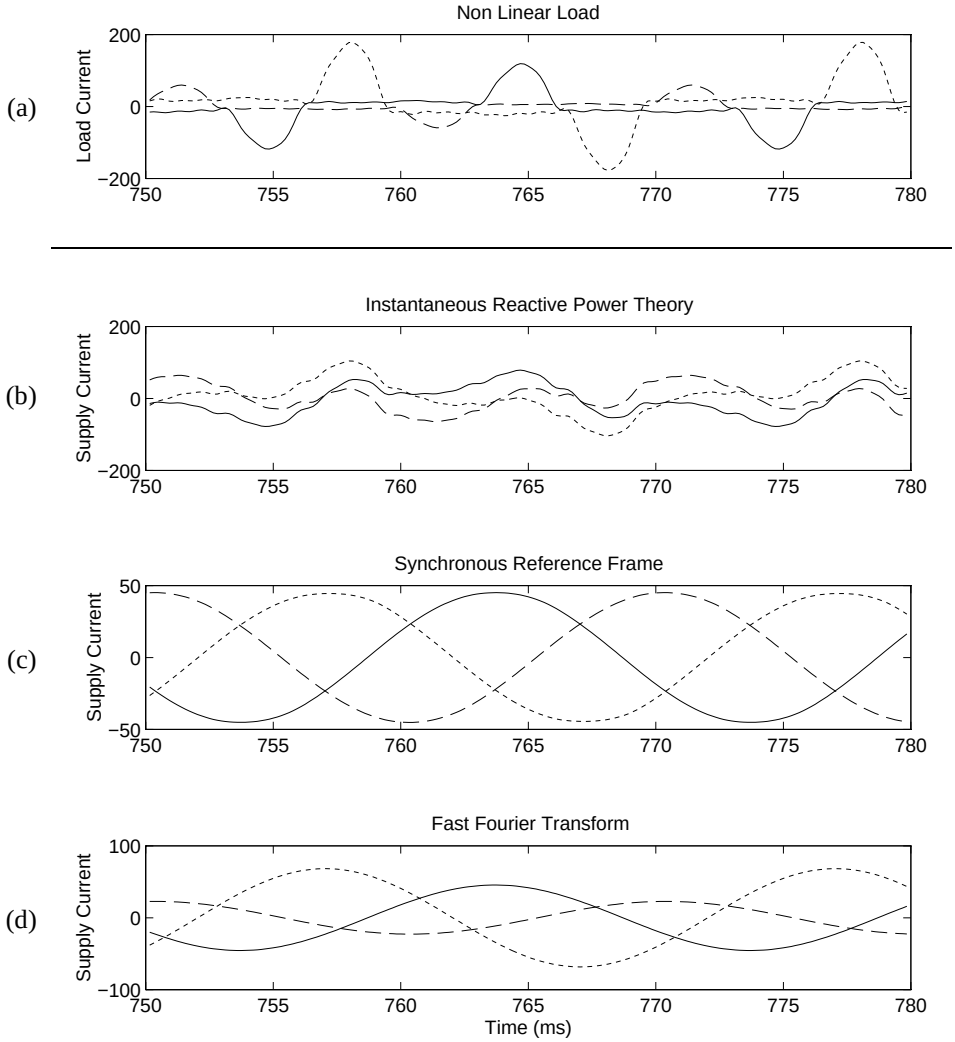

Figure 3.5 Comparison of steady-state currents for three harmonic isolation methods with an unbalanced load.

Errors introduced by the imbalance in the SRF method lead to unwanted load balancing and would require an inverter with a far larger power rating. The IRPT method has a large third harmonic component left in the supply current. Further simulation has shown that if the load is balanced and no triplen harmonics are present this method provides very good compensation with an excellent transient response. The FFT method correctly compensated the three phases and preserved the relative magnitudes of the currents. A similar result holds for Notch Filtering and Sinusoidal Subtraction.

3.4.1 Instantaneous Reactive Power Theory

An IRPT based controller was used with four different load current characteristics and the resulting supply currents are shown in Figure 3.6. The four load currents were:

(a) balanced load with no triplen harmonics

(b) unbalanced load with no triplen harmonics

(c) balanced load with triplen harmonics

(d) unbalanced load with triplen harmonics

When the load current is balanced and no triplen harmonics are present the IRPT method compensates for harmonics very well and the supply currents are sinusoids. If however, the current drawn by the non-linear load is unbalanced, the true harmonic current is not isolated from the load current and this results in residual harmonics in the supply current.

Triplen harmonics are ignored by an IRPT based active filter controller. As a result, the supply current will have triplen harmonics present and this is apparent from Figure 3.6(c). The worst situation is where an unbalanced load current contains triplen harmonics. Unfortunately this type of load is commonly caused by fluorescent lighting and small electrical appliances using switching power supplies, particularly personal computers and televisions. Results given in Figure 3.6(d) show that IRPT is not particularly good at compensating this type of load current.

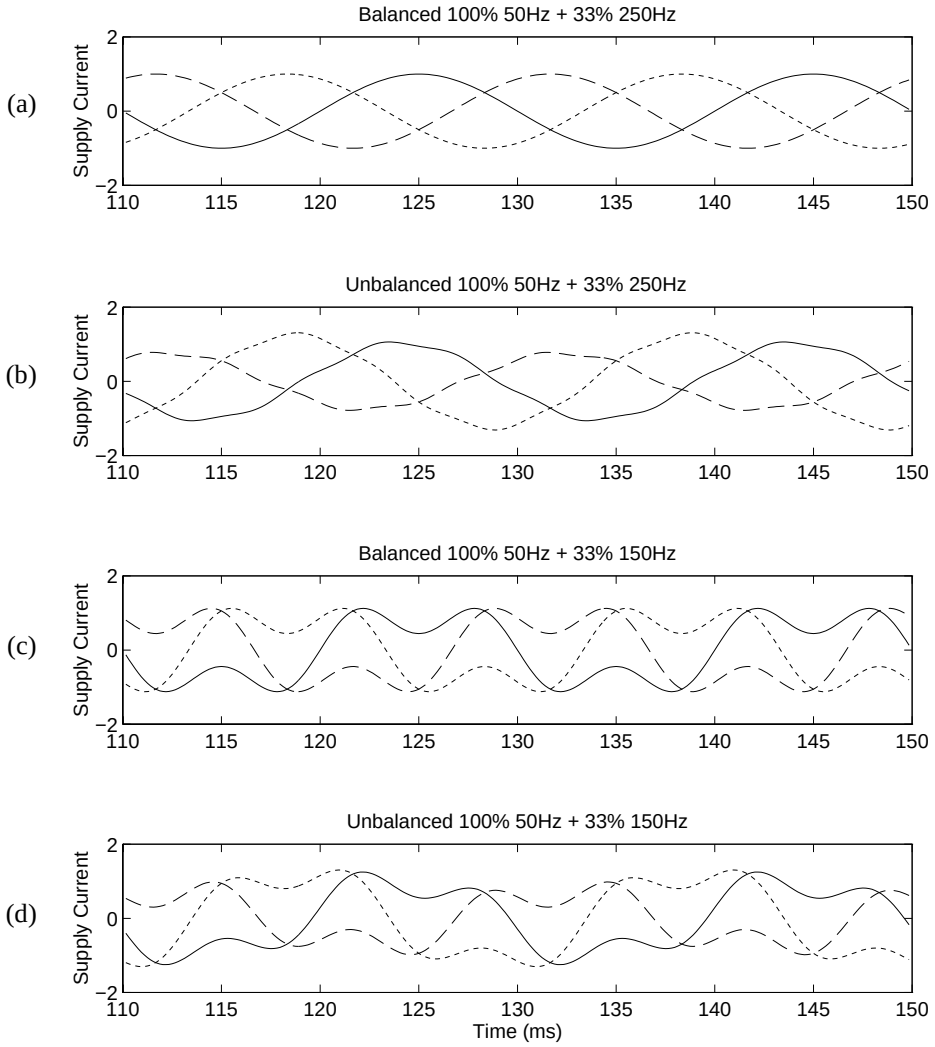

Figure 3.6 The resulting supply currents for the four possible load currents when IRPT is used to isolate the harmonic currents.

3.4.2 Synchronous Reference Frame and Fast Fourier Transform

An unbalanced current was simulated using Simulink and was used as the input to the Synchronous Reference Frame and the Fast Fourier Transform harmonic isolation methods and is shown in Figure 3.7(a). Figure 3.7(c) shows that the resulting supply current for the SRF harmonic isolation method remains balanced when the load current is unbalanced. The FFT technique, being typical of all single phase harmonic isolation techniques, gives an unbalanced supply current, as shown in Figure 3.7(e).

The compensating currents are quite different for each technique. In Figure 3.7(d), the compensating currents produced by the FFT technique are all the same waveshape. The waveshapes are different for the SRF generated compensating current in Figure 3.7(b). This

compensating current performs load balancing by introducing a fundamental component to the compensating current. The three compensating currents in Figure 3.7(d) have no fundamental component and so the supply currents in Figure 3.7(e) remain unbalanced in the same proportions as the load currents. This is different to the supply current for the SRF method, which has been averaged across the three phases, as shown in Figure 3.7(c). Load balancing increases the current through the inverter, leading to greater losses in the semiconductors, and increases the overall operating expense of the active filter. A bigger inverter will be needed to supply the larger currents and will increase the capital cost of the active filter.

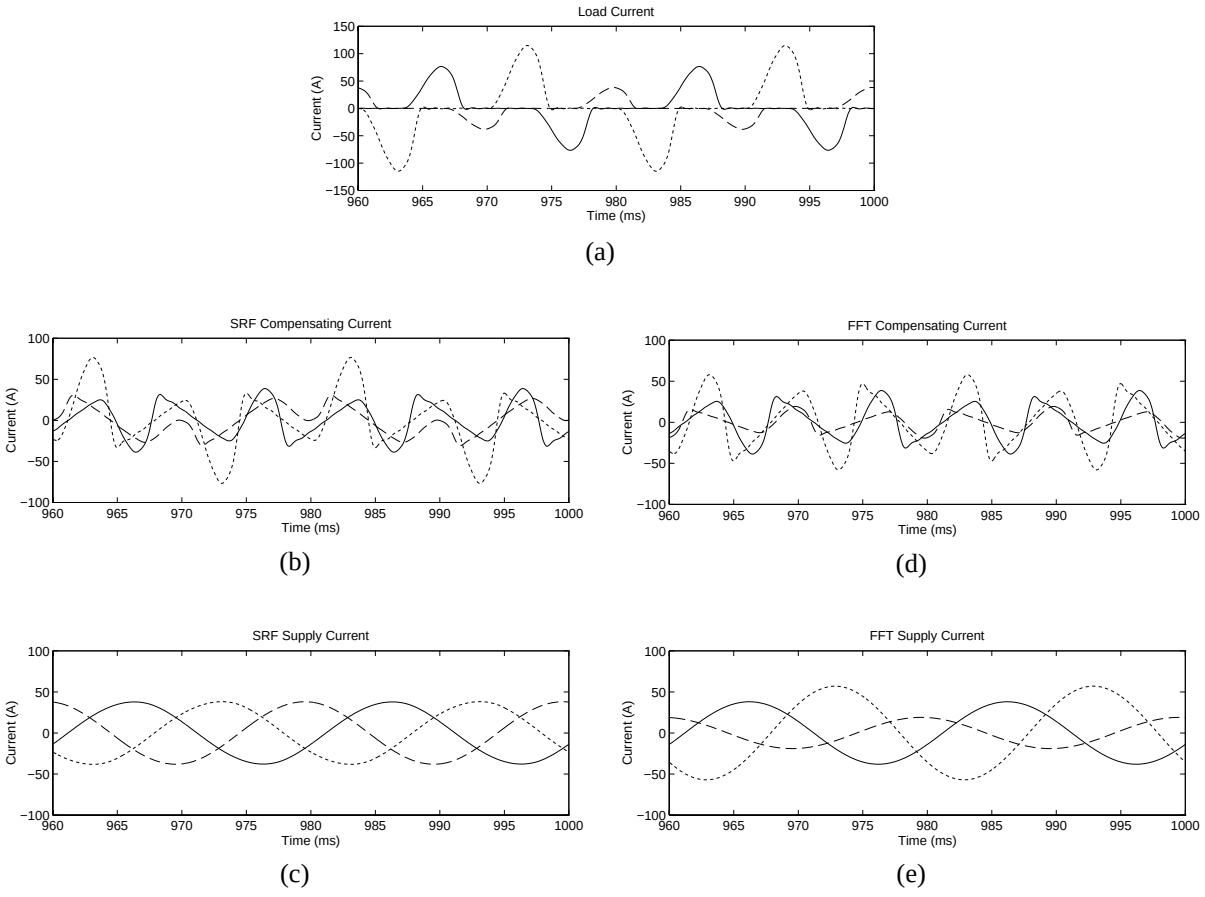

Figure 3.7 Load current (a) and supply currents for unbalanced operation of the SRF method (b)-(c) and the FFT method (d)-(e).

3.5 Summary

Industrial and commercial loads are typically unbalanced and contain harmonic currents. Five methods of determining the compensating current for an active filter

compensating this type of load have been evaluated with Simulink and MATLAB. Single phase techniques (Notch Filtering, Sinusoidal Subtraction and Fast Fourier Transform) provide good steady-state compensation but the three phase techniques (IRPT and SRF) do not correctly compensate current harmonics in unbalanced three phase systems.

The IRPT method is unable to compensate triplen harmonics and provides unsatisfactorily performance. The SRF method assumes that loads are balanced and is unable to distinguish between load imbalances and third harmonic components. This results in the SRF method performing undesired load balancing and would require a larger inverter than for the three single phase methods.

Of the single-phase harmonic isolation methods, the Fast Fourier Transform (FFT) based method gives the most accurate compensation, bringing the THD of the resulting supply current to 0.0%. Unfortunately, the two cycle delay in generating the correct compensating currents does affect the transient performance of this method. Sinusoidal subtraction almost completely compensates for the load current and responds more rapidly to transients than the FFT method. The simplest technique is Notch Filtering, but this has the disadvantage of having the slowest transient response. The next chapter discusses the hardware used to test the single-phase harmonic isolation techniques implemented with a Digital Signal Processor (DSP).

Chapter 4
Hardware Implementation

4.1 Experimental Apparatus

The work described in this thesis evaluated the different harmonic isolation methods through simulation (as discussed in the previous chapter) and by experimentation. This chapter describes the experimental apparatus, both 'off the shelf' equipment and that designed specifically for this work.

A 32 bit floating point Digital Signal Processor (DSP), the TMS320C30 from Texas Instruments, formed the core of the equipment. An IBM compatible personal computer provided electrical power for, and control of, this DSP. Data acquisition was performed by an eight channel, 12 bit, analogue to digital converter (ADC). The sampling of the ADC was controlled by a Sample Rate Multiplier (SRM) that synchronised the active filter to the mains frequency. Compensating current signals were transmitted via a Field Programmable Gate Array (FPGA) to a power inverter or power system model through a high speed synchronous serial link. The FPGA generated the transistor switching signals from the compensating current signal.

Figure 4.1 is a photograph of the equipment used in the experimental evaluation of harmonic isolation techniques. To the left of the photograph is the Tektronix TDS-510 digital oscilloscope used for making the waveform captures used in this thesis. A Hewlett-Packard Dynamic Signal Analyser performed the frequency and filter analysis for this work. Two personal computer screens can be seen on the bench. The left most is for the computer that contains the Texas Instruments XDS-510 debugger (with more detail given in §4.2.3) and the right screen shows the PC software used to control the DSP. The circuit boards that make up the hardware implementation of the three phase active filter are to the left of the debugger's monitor.

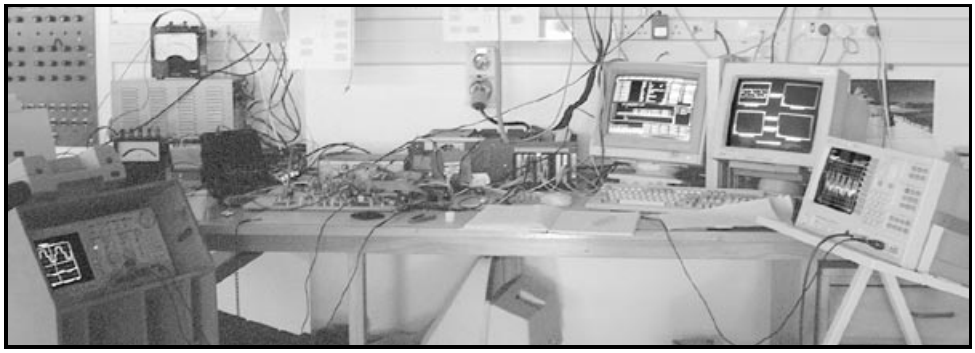

Figure 4.1 Photograph of the experimental apparatus used to evaluate the performance of harmonic isolation methods.

Figure 4.2 shows this equipment in diagrammatic form, with each block discussed in more detail in following sections.

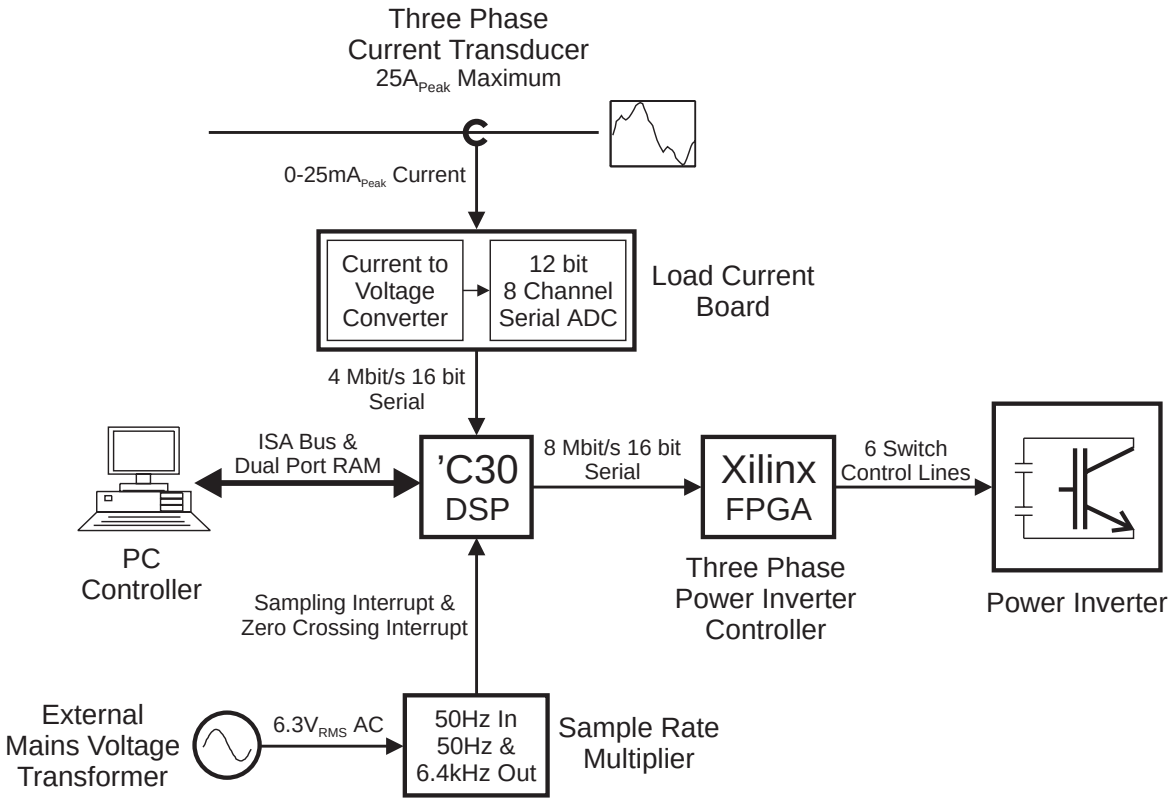

Figure 4.2 External connections of the equipment used for evaluating harmonic isolation techniques.

4.2 Digital Signal Processor

4.2.1 Sonitech Spirit-30 DSP Development Board

The TMS320C30 (also referred to as the ’C30) DSP was part of a Sonitech Spirit-30 development board that was mounted in an ISA slot in an IBM compatible personal computer. The DSP development board provided a convenient platform from which to develop DSP based applications and provided power for the DSP. The development board had external connectors for the two synchronous serial ports These peripherals will be discussed in more detail in the next section.

A software development library provided by Sonitech allowed a program written in Borland C or Microsoft C to communicate with programs running on the DSP. This software library was responsible for loading the DSP programs into the DSP’s high speed dual-port static RAM and will be discussed in greater detail in §6.4.1. Rudimentary debugging tools were provided by Sonitech, but were superseded by the XDS-510 Debugger (discussed in §4.2.3).

It is intended for the DSP to control a real active filter, and so the two serial ports were used for all external connections. These are much easier to optically isolate than parallel interfaces due to the reduced number of signal lines.

4.2.2 Texas Instruments TMS320C30

The TMS320C30 is a 32 bit floating point DSP and was originally developed in 1987 (Lin, Frantz and Simar, 1987). Many on-chip peripherals are provided to enhance its operation. Figure 4.3 shows the block diagram of the ’C30 from the *C3x User’s Guide* (Texas Instruments, 1994).

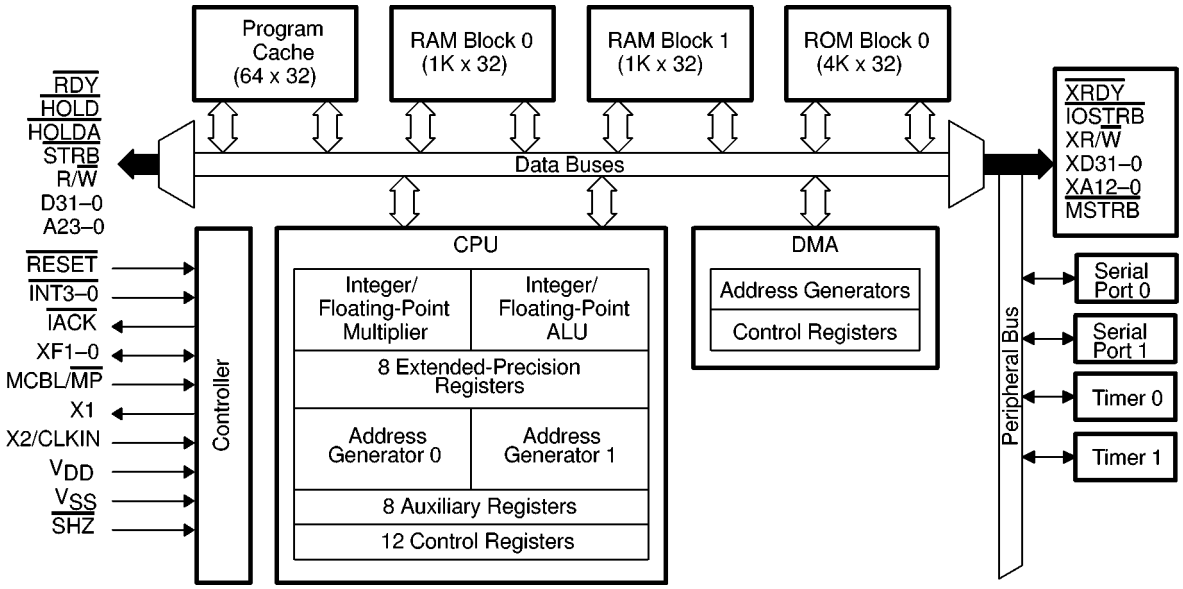

Figure 4.3 Block diagram of the TMS320C30 DSP, from the *C3x User's Guide* (Texas Instruments, 1994).

The DSP has two banks of 1k × 32 bit high speed Static RAM on the silicon die that can be accessed in a single clock cycle. This 'internal RAM' should be used for stack storage and for calculations that are speed critical. A 64 word program cache speeds up the execution of program sections that are stored in external dual port RAM (on the Spirit-30 development board). Two high speed (eight million bits per second) synchronous serial ports are provided for communicating with external devices or another 'C30 DSP. Two 32 bit timer/counters are provided but were not used in this implementation of an active filter controller. A Direct Memory Access (DMA) controller on the chip allows high speed memory transfers. This has the ability to move large pieces of data between internal and external memory and the on-chip peripherals without using the DSP's CPU. This CPU provides integer and floating point calculation units and a flexible set of registers; eight 40 bit extended registers and eight 32 bit auxiliary registers.

An important feature of DSPs are separate memory busses (Lin *et al.*, 1987) which allow the CPU and the DMA Controller to access memory at the same time. The Primary and Expansion busses on the 'C30 allow simultaneous memory access to peripherals (such as analogue to digital converters) and data memory (either external or internal RAM). The dedicated 32 bit floating point multiplier performs multiplications in a single clock cycle, giving the DSP a significant speed advantage over conventional microprocessors and microcontrollers. The DSP that was used operated with a 33MHz clock, giving a processor

cycle length of 60ns. Later revisions of the processor are capable of operating at 50MHz, giving a 40ns processor cycle length (Texas Instruments, 1994).

4.2.3 Texas Instruments XDS-510 In-Circuit Debugger

Towards the completion of experimental work an in-circuit debugger was provided by Texas Instruments. This connected to a dedicated debugging port on the DSP and allowed much more comprehensive debugging than that provided by Sonitech with the Spirit-30 development board. The debugger was installed in another IBM compatible personal computer which allowed debugging of communications between the DSP and the host PC, something the Sonitech debugger could not perform.

This debugger supported DSP programs written in C and allowed single stepping through the source code, rather than the resulting assembly language instructions. Precise control of the DSP was obtained by single stepping and setting breakpoints. One of the most useful features provided by the XDS-510 was instruction cycle counting. The debugger was able to count the number of clock cycles required to execute a section of code. This was then used to measure the computational complexity of each harmonic isolation technique, the results of which are discussed in more detail in §8.2.

4.2.4 Synchronous Serial Ports

All connections to the DSP were made through its two synchronous serial ports, SP0 and SP1. These came out to ten pin wafer headers on the Spirit-30 development board with six signal lines and two grounds supplied, leaving two spare pins on each. On SP0, one of the spare pins supplied power to a serial port buffer used to extend the cable distance, while the other pin connected to the INT3 input on the DSP (one of four external interrupts). Spare pins on the SP1 header were connected to the INT2 and INT1 inputs, for the Zero Crossing Detect interrupt and the Sampling interrupt respectively.

Six data lines were connected to each serial port; three for transmit and three for receive. Each serial port has the following signals:

CLKX: Transmit Clock

CLKR: Receive Clock

FSX: Transmit Frame Synchronisation

FSR: Receive Frame Synchronisation

DX: Transmit Data

DR: Receive Data.

The two halves of the serial port are independent and may operate with completely different configurations. This is taken advantage of in the interfacing of an ADC to the DSP and is discussed in more detail in §4.4.3. The clock and frame synchronisation pins can be configured as inputs or outputs and all pins have configurable polarity (active high or active low). This gives great flexibility but does make configuration more complicated. The serial ports are double buffered; this allows the programmer to supply a new value or read a value into the Data Transmit Register (DXR) or the Data Receive Register (DRR) while more data is being shifted into the Transmit Shift Register (XSR) or out of the Receive Shift Register (RSR). Interrupts can be generated by the serial ports when data is transmitted or received, removing the need to poll the serial port before sending or receiving new data. The DMA controller may also be used transfer large blocks of memory to and from the serial ports without CPU intervention.

4.3 Mains Frequency Synchronisation

Standards dealing with current harmonics provide limits up the fiftieth harmonic (AS 2279.1, 1991; IEC 1000-4-7, 1991; IEEE Std. 519, 1992; BS EN 61000-3-2, 1995) and therefore the active filter must be able to measure up to at least the fiftieth harmonic, even though the inverter may not compensate frequencies this high. The frequency spacing of the FFT output is the inverse of the sample period (Poularikas *et al.*, 1993), as given by Eqn. (4.1).

$$\Delta f = \frac{1}{T} \tag{4.1}$$

If the sampling period is one mains cycle, Δf will be the mains frequency and so each output point in the frequency domain will be a harmonic (see §2.2.5 for more detail on the operation of the FFT). To resolve the fiftieth harmonic, $N/2$ must be 50 or greater, and therefore N will be at least 100. Additionally, the FFT routines used require that the number of samples is an integer power of two (Tessarolo, 1991). For this work 128 samples per mains cycle was chosen as this satisfies the above criteria while keeping the processing load of the DSP to a minimum.

The digital filters and Fast Fourier Transforms used in the experimental work are written with the assumption that there are 128 samples per mains cycle, giving a nominal sampling frequency of 6.4kHz. The mains frequency is never exactly 50Hz and changes

slowly over time as the power system responds to changing loads. The Fast Fourier Transform (FFT) assumes there are an integer number of samples in each time 'window' (Brigham, 1974) and so the active filter must accommodate changes in mains frequency.

4.3.1 Frequency Multiplication

A Sample Rate Multiplier (SRM) was used to multiply the mains frequency by 128 to give exactly 128 samples in each mains cycle. This frequency multiplication was performed by a CMOS Phase Locked Loop (PLL), the Philips HEF4046, and a HEF4040 12 bit binary ripple counter acting as a divider. The block diagram of the frequency multiplier is shown in Figure 4.4.

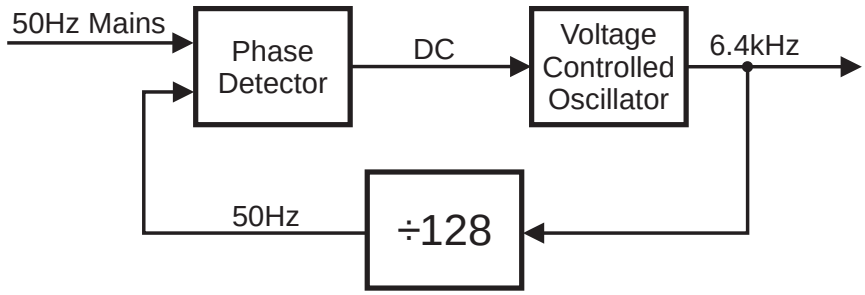

Figure 4.4 Block diagram of the Sample Rate Multiplier using a Phase Locked Loop.

The input to the SRM was a clipped representation of the 50Hz mains voltage and is a square wave with exactly the same frequency as the mains voltage. A Voltage Controller Oscillator (VCO) generated a square wave with a frequency ranging from 6.27kHz to 6.53kHz. This was divided down by 128 to give a square wave with a frequency ranging from 49Hz to 51Hz. The phase of the mains signal was then compared to the divider's output with a Phase Detector in the HEF4046 PLL. The output of the Phase Detector was filtered to give a DC voltage that controlled the VCO output frequency. A voltage of $0.9V_{DC}$ gave the lower frequency limit (6272Hz) and $4.1V_{DC}$ gave the upper limit (6528Hz).

The filtered DC voltage changed so that the divided down square wave was exactly in phase with and the same frequency as the mains voltage. Once this occurred the SRM was said to be "in lock". The "capture range" is the range of input frequencies that the SRM was able to lock onto, in this case 49Hz to 51Hz. This range was kept small so that any noise present on the VCO input gave a small change in the frequency of the output. If the mains frequency is less than 49Hz the VCO output will remain at the lower limit of 6272Hz.

Similarly, if the mains frequency exceeds 51Hz, the VCO frequency will not exceed the upper limit of 6528Hz. When the PLL input is outside the capture range the sampling will not be synchronous and the active filtering will be erroneous. The time taken for the SRM to lock once power was applied averaged ten seconds. When the capture range was increased the SRM locked faster, but was more susceptible to noise at the VCO input.

4.3.2 Zero Crossing Detection

An isolated Zero Crossing Detector (ZCD) was used to generate a square wave of the same frequency as the mains voltage. This signal was derived from a $6.3V_{AC}$ transformer which powered the "live side" of the SRM. The ZCD signal was fed into the DSP to provide an absolute phase reference and was used by the PLL that performed the frequency multiplication.

4.3.3 DSP Interrupts

The SRM produced a square wave with a frequency 128 times that of the power system's fundamental frequency. This was used to generate sampling interrupts for the DSP by converting each rising edge into a negative going pulse, as shown in Figure 4.5. This was necessary as the TM320C30 DSP used level triggered interrupts rather than edge triggered interrupts. If the interrupt line is held low for more than three clock cycles (180ns) multiple interrupts can be triggered (Texas Instruments, 1994). The Schmidt triggered NAND gates and an RC timing circuit generated 100ns wide pulses on each rising edge of the input. The SRM output and the ZCD signal were both turned into DSP interrupts.

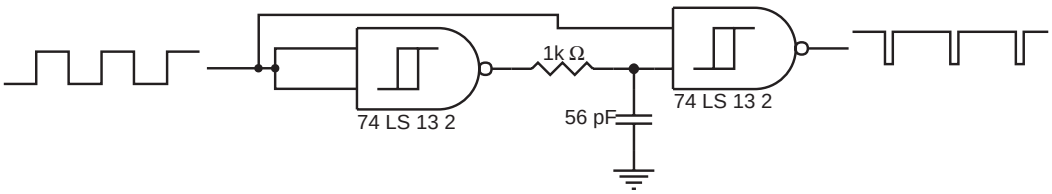

Figure 4.5 Interrupt generation circuitry for the TMS320C30.

4.3.4 Implementation of the Sample Rate Multiplier

The full schematic diagram for the Sample Rate Multiplier is given in Appendix A.3 and a photograph of the assembled SRM printed circuit board shown in Figure 4.6. The left

hand side of the board is the isolated Zero Crossing Detector. An opto-isolator passed the ZCD signal to the rest of the circuit. The frequency multiplier was implemented with the upper two integrated circuits (ICs) on the right hand side, the HEF406 and HEF4040. A Schmidt triggered NAND gate shown in the bottom right corner produces the interrupt pulses for the DSP. These are transmitted by the two core cables to the Load Current board, that then sends the signals onto the Serial Port connectors of the Spirit-30 development board.

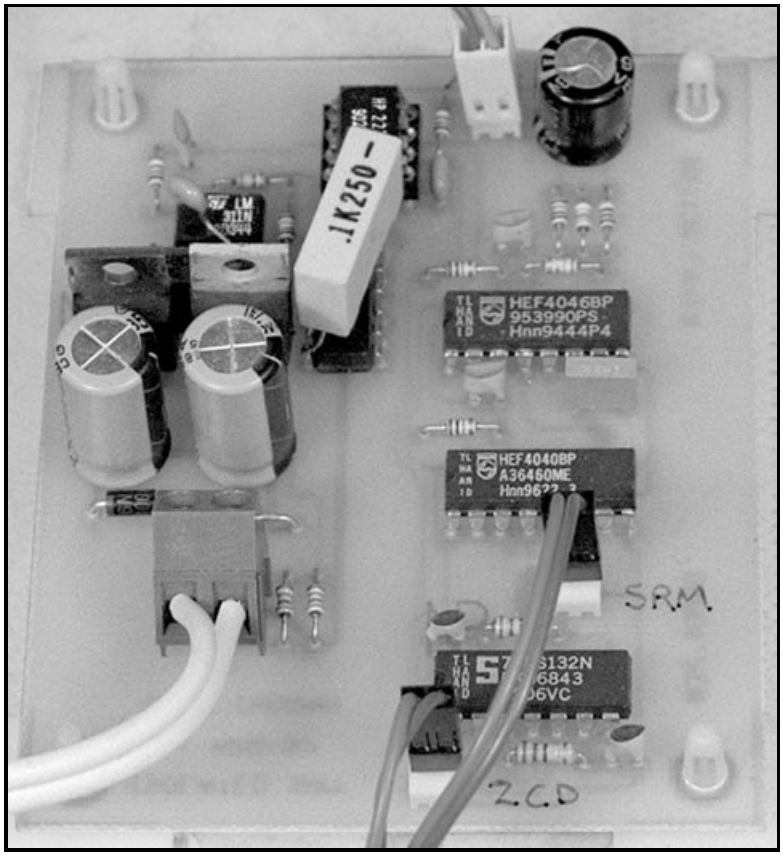

Figure 4.6 Photograph of the Sample Rate Multiplier printed circuit board.

4.4 Load Current Measurement

The digital Active Filter controller required knowledge of the current drawn by the harmonic producing load for it to generate the compensating current. The load current is an analogue quantity and so had to be sampled and digitised for use by the DSP. A block diagram of the load current data acquisition system is shown in Figure 4.7. Each of the blocks in the figure will be discussed in the following sections.

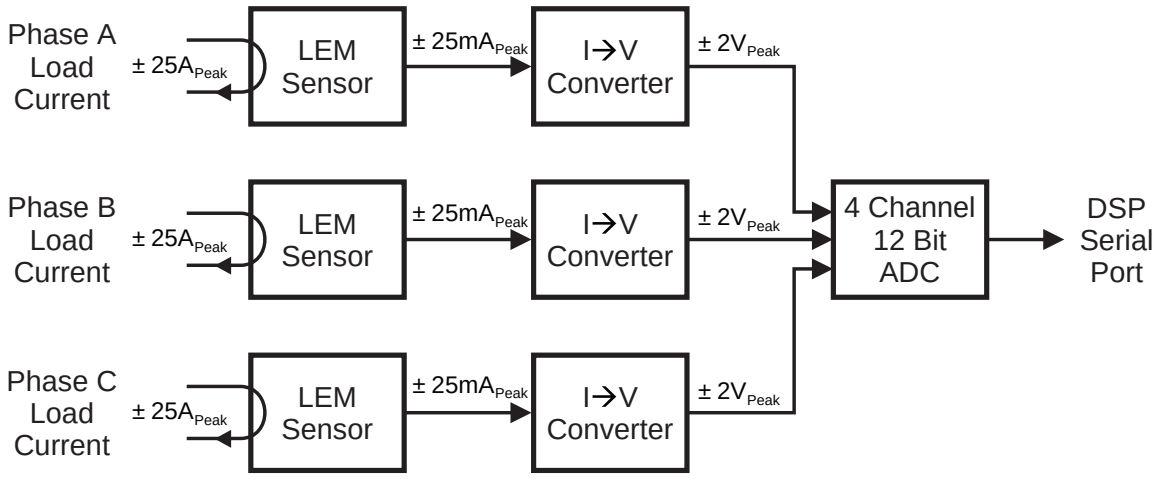

Figure 4.7 Load current measurement and acquisition system.

4.4.1 Load Current Sensors

The load current for each phase was passed through a 'LEM Sensor' solid state current transducer. The LEM sensor is an active current nulling device rather than a current transformer (CT). Magnetic flux generated by the load current is matched by an active circuit so the net flux (determined internally by the sensor) is zero. This allows the sensor to accurately measure DC and high frequency currents. The maximum current that could be measured accurately by the sensors was $\pm 25A_{Peak}$ when a 1000:1 'turns ratio' was selected. The output of the LEM sensor is a current, with a maximum range of $\pm 25mA_{Peak}$. The signal from the LEM sensors to the load current circuit board was kept as a current to reduce susceptibility to interference by stray electric fields.

4.4.2 Current to Voltage Converter

The Maxim MAX186 analogue to digital converter (ADC) required the current from the LEM sensor be converted to a voltage of the appropriate input range. The simplest current to voltage converter is the resistor which has the disadvantage of presenting a non zero impedance to the source of the input current. This can produce errors if the device producing the current has very little drive (a low voltage output) or if it does not produce a constant current as its output voltage changes (Horowitz and Hill, 1989). A transimpedance amplifier was used to perform the current to voltage transformation while providing a near zero input impedance for the LEM sensor. The circuit diagram of the transimpedance amplifier is given

in Figure 4.8. Offset and gain adjustments are provided to compensate for the variations in the LEM sensor, op-amp and ADC.

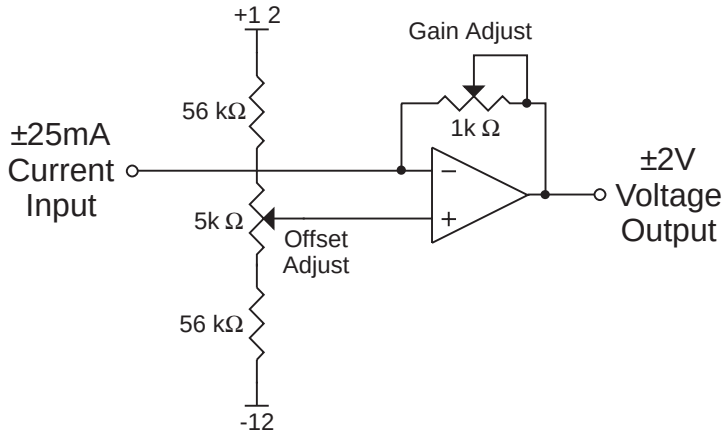

Figure 4.8 Transimpedance amplifier for interfacing a LEM current sensor to an ADC input.

4.4.3 Analogue to Digital Converter

A current resolution of 0.1A over a $\pm 200A_{Peak}$ range was deemed adequate for a full sized active filter that breaks the full current range into 4000 points. A 12 bit ADC resolves its input range into 4096 points and so is suitable for this data acquisition task.

The MAX186 is a 12 bit, four channel, serial ADC and was used to sample the output of the three transimpedance amplifiers. This ADC may also operate as an eight channel single ended unipolar (0V to +4V) converter. The MAX186 had an internal voltage reference which removed the need for an external precision reference. Internal details of the MAX186 are shown by the block diagram in Figure 4.9.

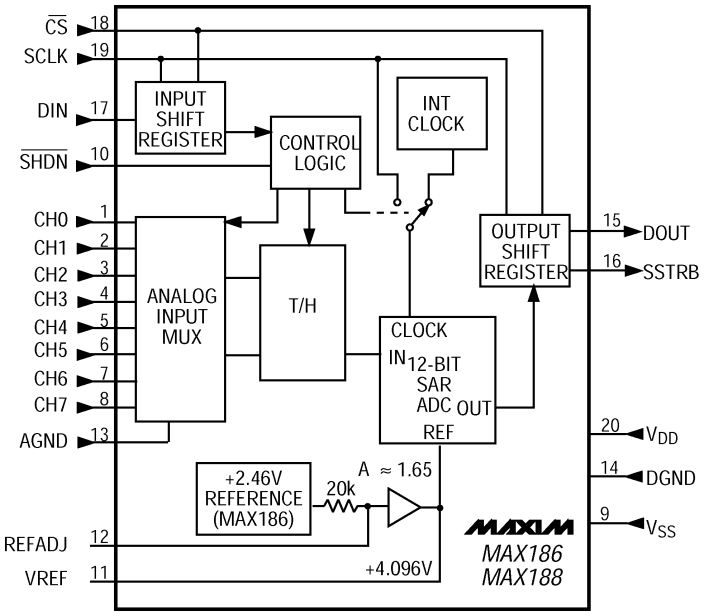

Figure 4.9 Internal block diagram of the MAX186 analogue to digital converter, from the MAX186 datasheet (Maxim, 1995).

The inputs to the ADC are multiplexed and therefore only one channel can be read at any one time, leading to slight phase errors between readings. The time delay between readings is not constant due to instruction caching and pipelining (executing several instructions in parallel) (Texas Instruments, 1994) and may introduce errors into the calculations. The observed results show these phase errors to be small and almost constant and were therefore neglected. If the errors were significant simultaneous sampling would avoid the problem and could be performed by three separate ADCs or a single ADC designed specifically designed for the task.

The DSP sent a ‘control byte’ to the ADC whenever a sample was triggered by an interrupt from the Sample Rate Multiplier. The control byte selected the following: the channel to be converted, unipolar or bipolar mode, and single ended or differential mode. Power saving options were provided but were not used in this application. The timing diagram for the interface between the MAX186 and the ’C30 shows the format of the control word and is presented in Figure 4.10.

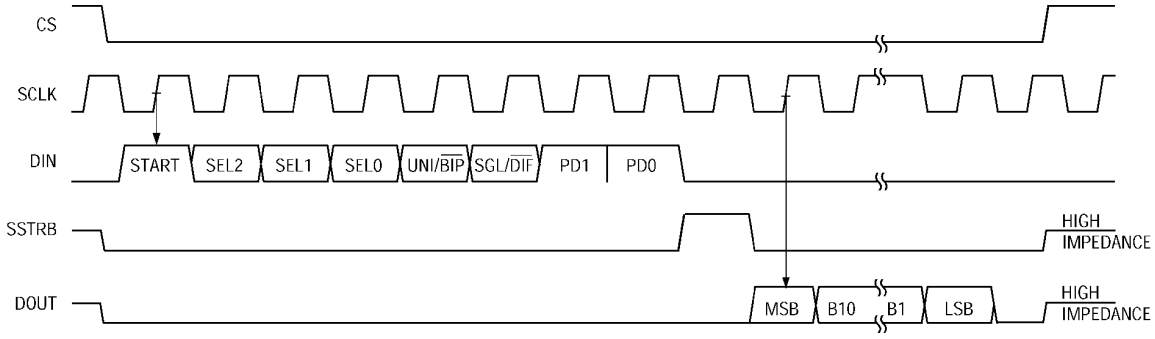

Figure 4.10 Serial communications format for the MAX186 ADC from the MAX186 Datasheet (Maxim, 1995).

A control byte is sent by the DSP to the ADC. Once the requested channel number has been shifted into the ADC conversion starts. When this is complete, the SSTRB line pulses to signify to the DSP that the converted sample is about to be transmitted. The most significant bit is sent first, then the remaining eleven bits of the reading, followed by four bits of padding to make the transfer sixteen bits long.

The recommended connection (Maxim, 1995) was not used as this required a connection to the XF0 pin which was not present on the DSP's serial port. An alternative was designed, with the connections shown in Figure 4.11. The Transmit Clock (CLKX) pin on the DSP drives the DSP's Receive Clock (CLKR) and the ADC's Serial Clock (SCLK) at 2MHz. Data transmissions from the DSP are set to 32 bits with Transmit Frame Sync (FSX) active low for the entire transmission. This generates a Chip Select ($\overline{CS}$) for the ADC that is longer than the 25 clocks required to transmit the result to the DSP. The DSP's Transmit Data (DX) pin sends the control byte to the Data In (DIN) pin of the ADC. One clock cycle after the pulse is generated from the ADC's Serial Strobe (SSTRB) line the Data Out (DOUT) pin on the MAX186 transmits the conversion result to the Receive Data (DR) pin on the DSP's serial port.

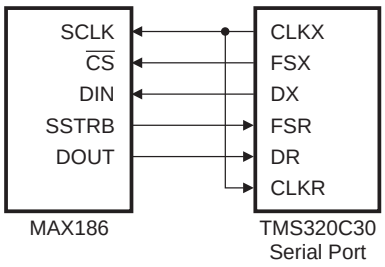

Figure 4.11 Connection diagram for interfacing the MAX186 ADC to the TMS320C30 DSP.

4.4.4 Load Current Board Implementation

The circuit schematic for the Load Current data acquisition board is given in Appendix A.1 and a photograph of the assembled board is shown in Figure 4.12.

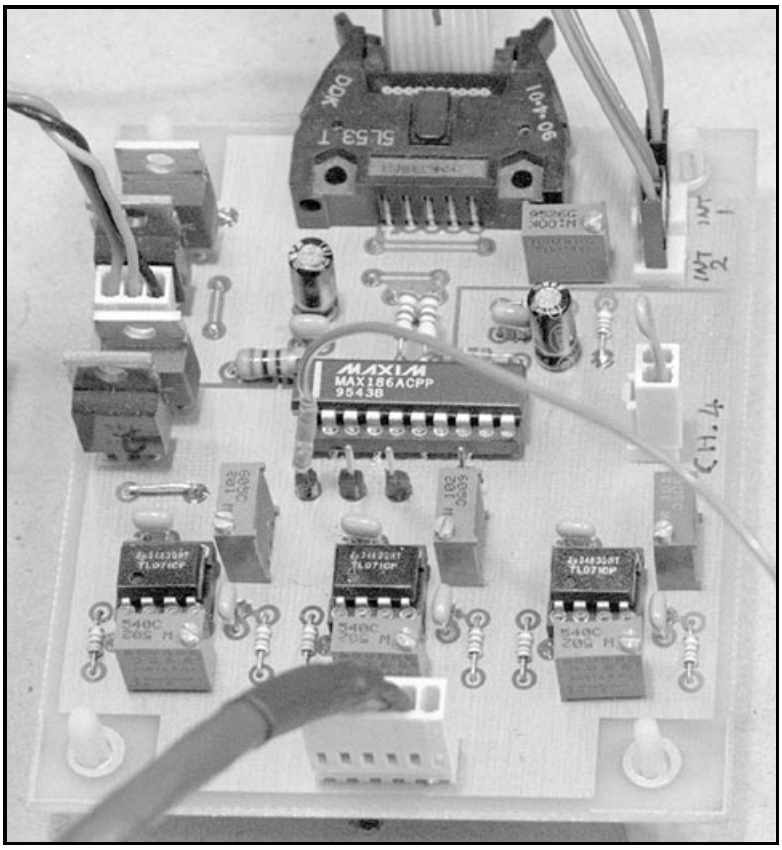

Figure 4.12 Load Current data acquisition board with three current to voltage converters and a four channel analogue to digital converter.

The board has separate analogue and digital power supplies to minimise noise conducted into the transimpedance amplifiers and ADC inputs by the digital section of the ADC. These were created with separate analogue and digital +5V and -5V regulators fed from the same ±12V supply. Separate analogue and digital ground planes reduced the effects of noise caused by ground return. Three phase currents from the LEM sensors enter through the header at the bottom of the board. The three transimpedance amplifiers with gain and offset adjustments, visible at the bottom of the photograph, fed into the centrally located MAX186. Serial transmissions to and from the DSP enter through the ten way IDC header at the top of the photograph. Two pin headers located next to the header connect to the Sample Rate

Multiplier board and bring the Zero Crossing Detect and Sampling interrupt signals to the external interrupt pins on the Spirit-30 serial port.

A header mounted to the right of the MAX186 connects to the fourth channel of the ADC used to sample the resulting load current in the experiments discussed in Chapter 8. This ‘supply current’ was created by summing the compensating current produced by a Digital to Analogue Converter (DAC) with the Load Current signal. The PC software displayed this waveform along with the mathematical addition of these quantities.

4.5 Switching Power Inverter

As part of the hardware implementation of the active filter a three phase voltage source inverter (VSI) was built. The rating of switches in the inverter were 15A and 600V, although these extremes were unable to be tested. The block diagram of the inverter is shown in Figure 4.13. Letters at the left and right hand sides of the IGBT Bridge represent the similarly named internal connections. Each of the blocks will be discussed in more detail in following sections.

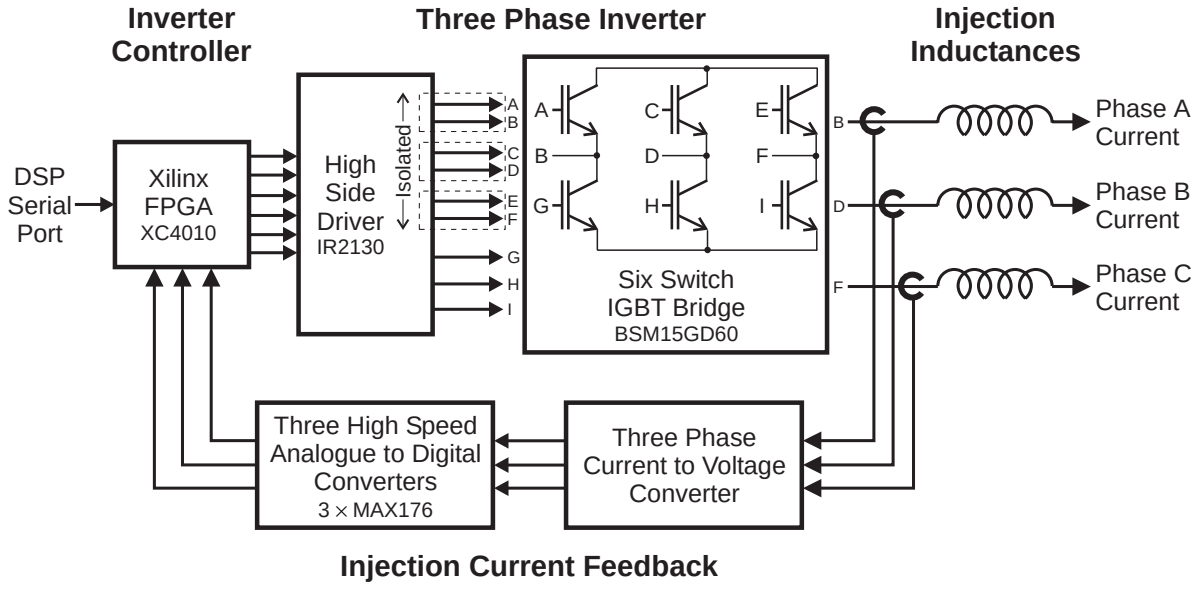

Figure 4.13 Block diagram of the three phase inverter used as part of the experimental work.

4.5.1 Inverter Controller

The inverter controller was implemented in a Xilinx FPGA that interfaced to one of the DSP's serial ports (the Load Current acquisition board connected to the other). The output of the FPGA was six switch control signals. Dead time protection was incorporated into the Inverter Controller to ensure that the upper and lower IGBT in each leg would never be conducting at the same time, preventing the possibility of conduction faults. Further discussion regarding the design and implementation of the controller can be found in Chapter 5.

Appendix B gives schematic diagrams for the Xilinx logic. These schematics are hierarchical and work down through several levels. A hardware description language, ABEL, was used to describe the behaviour of a timing state machine for the dead time protection circuitry. The source code describes the behaviour of a state machine along with the circuits inputs and outputs and is given in Appendix B.6. This prevented one switch's control line from going active until the other switch's control signal had been low for at least thirteen clock cycles. The specified turn off delay, $t_{d(\text{off})}$, and the fall time, t_f, for the module used (discussed in the next section) combined to give a total turn off time of 800ns (Siemens, 1996). An 8MHz clock gave a deadtime of 1.64μs, providing a safety factor of two.

4.5.2 Three Phase Inverter

A "High Side Driver" from International Rectifier was used to provide isolated gate drives for each of the top switches in the inverter. The voltage reference for each of these is the output point for that phase, leading to large potential differences between the individual IGBT emitters. This driver IC removed the need for isolated power supplies and pulse transformers for each of the high side transistors and provided high current non-isolated gate pulses for the lower switches in each leg, ensuring the inverter switched as quickly as possible.

The six IGBT transistors (and their anti-parallel diodes that are not shown) were contained in a single module from Siemens (Siemens, 1996). This package had extremely good thermal characteristics and was bolted directly onto a single piece of heatsink. The voltage output from each leg was applied to an inductor. By switching the IGBTs appropriately, the current through the inverter could be made to follow a reference.

The IGBT module and high side driver were mounted on the same printed circuit board along with the power terminals. Figure 4.14 shows the assembled circuit board with the

heatsink mounted directly onto the IGBT module. The high side driver is visible at the back right corner of the board. DC terminals are mounted on the reverse (topside during normal operation) at the front, with the three phase output terminals on the left hand underside of the printed circuit board.

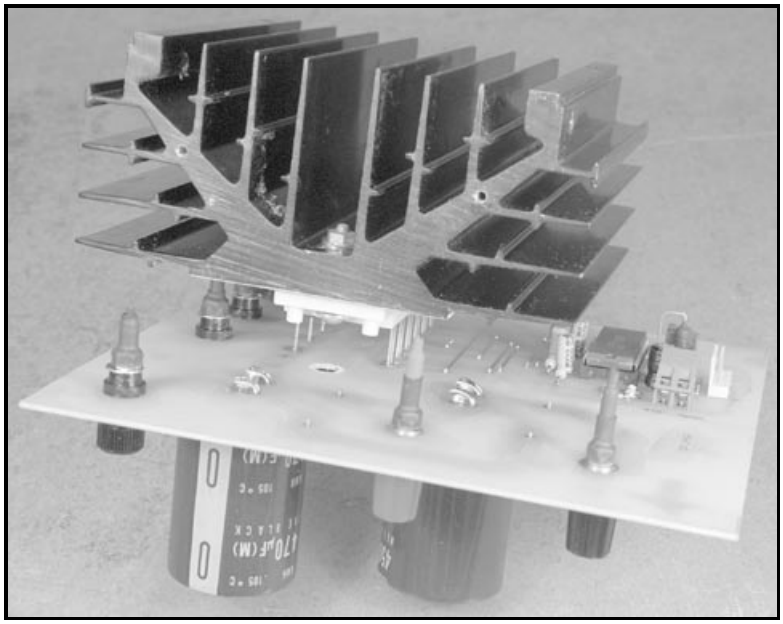

Figure 4.14 Underside of the three phase inverter built with a 15A 600V Siemens IGBT module and an International Rectifier high side driver.

4.5.3 Injection Current Feedback

Hysteresis current control (discussed in more detail in §5.2) regulated the current in the voltage source inverter. This control strategy uses current feedback and so the output current from the inverter was measured. The current transducers were the same type as used for measuring non-linear load currents, along with almost identical transimpedance amplifiers. The Inverter Controller is implemented digitally and so all hysteresis control decisions were made digitally. High speed sampling of the injected current was used in this decision making. It was found that a sampling frequency in excess of 50kHz was required to keep current ‘overshoots’ at acceptable levels, with faster sampling giving better results (Ingram and Round, 1997).

Three MAX176 high speed 12 bit serial ADCs were used to sample the injected current. Twelve bit ADCs were used to simplify the logic inside the Xilinx FPGA as all other

calculations used 12 bit signed integers. The sampling clock was generated by the FPGA and was 260kHz, the fastest that the ADC would accommodate. Figure 4.15 shows a photograph of the three injection current boards and Appendix A.2 gives the schematic diagram for each board.

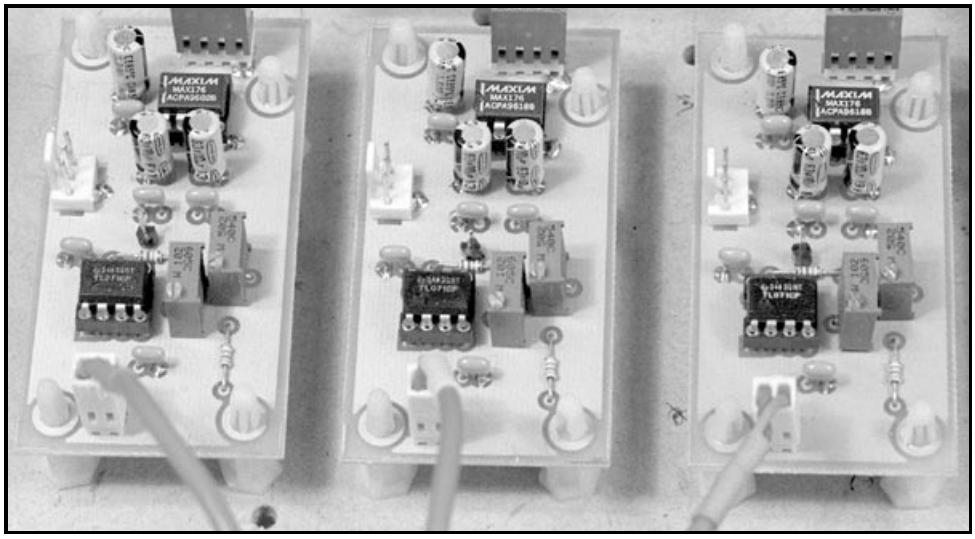

Figure 4.15 Photograph of the injection current feedback acquisition boards.

The transimpedance amplifier is visible at the bottom of each circuit board. As with the load current measurement board, gain and offset adjustments are provided to compensate for variations in the LEM current sensors. The MAX176 ADC is visible at the top of each board, along with the serial data connector that interfaces each data acquisition board to the Xilinx FPGA.

4.6 Summary

This chapter has described the hardware used to evaluate the harmonic isolation techniques in this work. ‘Off the shelf’ equipment has been interfaced to electronics built specifically for the task to provide a custom solution. All external processing of data by the active filter is as 12 bit signed integers with the DSP’s 32 bit floating point format used for all numeric processing and digital filtering calculations. A hysteresis current controller has been implemented in an FPGA that generated the switching signals from a digital reference produced by the DSP. This reduced the amount of analogue circuitry and enhanced the

system’s immunity to noise. The following chapter discusses the design of the inverter controller and presents performance results from simulation and experimental work.

Chapter 5
Digital Current Controller

5.1 Control Strategies

Active filters eliminate harmonic currents by injecting them back into the power system with a 180° phase shift. A controlled current inverter is required to generate this compensating current. Current Source Inverters (CSI) naturally have a current as their output and are well suited to active filtering applications (Hayashi and Takahasi, 1991; Fukuda and Endoh, 1995). An alternative is the Current Regulated Voltage Source Inverter (CR-VSI) which is specially controlled so that its output is a controlled current (Malesani and Tomasin, 1993). CSIs have superior controllability and reliability when compared to Voltage Source Inverters (VSI), but most research has focused on the use of VSI inverters because of their lower operating costs (Fukuda and Endoh, 1995). This is due to energy storage in capacitors being cheaper than in inductors because of lower operating losses.

Several techniques exist for controlling the current from VSI inverters. These fit into three main categories: hysteresis, ramp comparison and predictive control (Brod and Novotny, 1985; Malesani and Tomasin, 1993; Kazmierkowski and Dzieniakowski, 1994; Holtz, 1997). Hysteresis current controllers (also referred to as “Bang-Bang” controllers) compare the desired current to the line current with some hysteresis. Unfortunately the switching frequency does not remain constant (Lorenz, Lipo and Novotny, 1997) and sub-harmonic components can be generated (Holtz, 1997). Ramp comparison keeps the switching frequency constant or at least bounded within a known range. A Proportional-Integral (PI) controller endeavours to keep the line current the same as the desired current by changing the duty cycle of the inverter (Lorenz *et al.*, 1997). Predictive current controllers select the optimum voltage vector from the inverter to best move a current around a vector plane. The future path of the current vector is predicted and this is used in the decision making (Lorenz *et al.*, 1997). This technique keeps a constant switching frequency but has variable current ripple, the opposite of the hysteresis controller (Brod and Novotny, 1985).

Some progress has been made in the field of constant frequency hysteresis current control. This has been achieved by changing the hysteresis band amplitude within the output waveform period (Malesani *et al.*, 1991). Further work has led to a technique where by the

current is predicted and is used to control the hysteresis limits; this locks the switching frequency to a synchronising clock (Malesani *et al.*, 1996). Even though this technique uses digital control, the hysteresis part of the current controller is still implemented with analogue electronics.

The hysteresis current control method was selected for this work because it is one of the simplest to implement (Brod and Novotny, 1985), and provides the fastest and most stable control available (Malesani and Tomasin, 1993). Rather than using analogue electronics to implement the controller, a Field Programmable Gate Array (FGPA) was used. The design decisions and performance of this controller are discussed in detail in the remainder of this chapter.

5.2 Hysteresis Current Control

Hysteresis Current Control (HCC) controls the switches in an inverter asynchronously to ramp current up and down through an inductor so that it tracks a reference current signal. A hysteresis current controller is implemented with a closed loop control system and is shown in diagrammatic form in Figure 5.1. An error signal, $e(t)$, is used to control the switches in an inverter. This error is the difference between the desired current, $i_{ref}(t)$, and the current being injected by the inverter, $i_{actual}(t)$. When the error reaches an upper limit, the transistors are switched to force the current down. When the error reaches a lower limit the current is forced to increase.

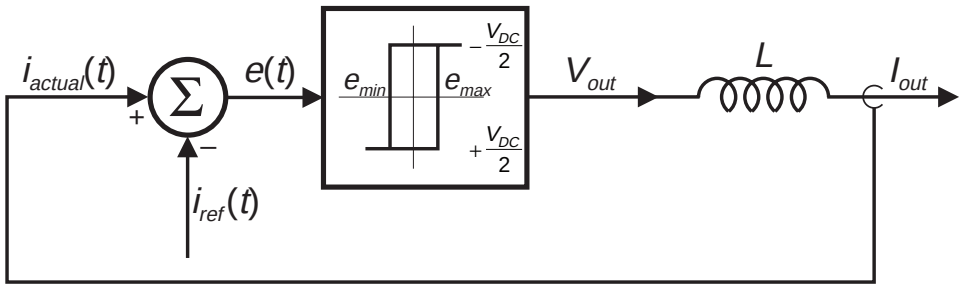

Figure 5.1 Hysteresis current control signal block.

The minimum and maximum values of the error signal are e_{min} and e_{max} respectively. The range of the error signal, $e_{max} - e_{min}$, directly controls the amount of ripple in the output current from the inverter and this is called the Hysteresis Band. The hysteresis limits, e_{min} and

e_{max}, relate directly to an offset from the reference signal and are referred to as the Lower Hysteresis Limit and the Upper Hysteresis Limit. The current is forced to stay within these limits even while the reference current is changing. The ramping of the current between the two limits is illustrated in Figure 5.2.

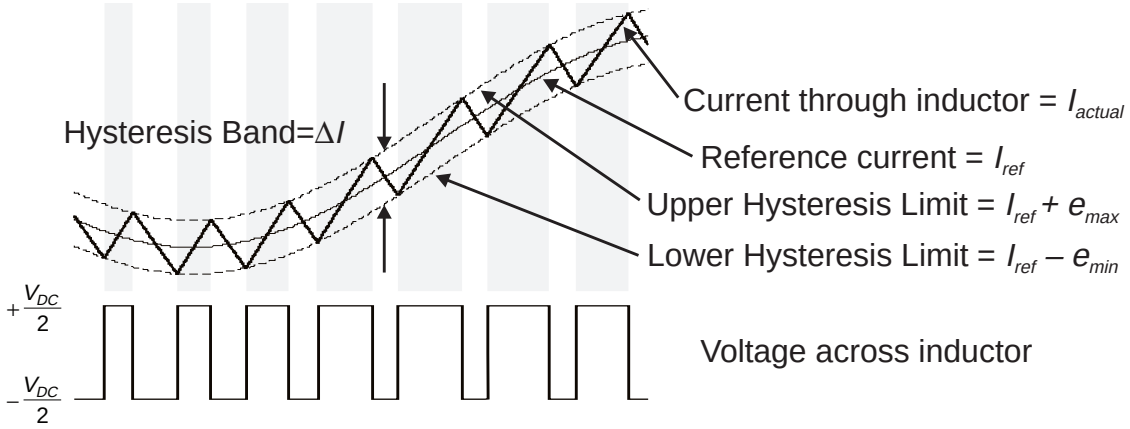

Figure 5.2 Hysteresis current control waveforms.

The switching frequency is altered by the width of the hysteresis band, the size of the inductor that the current flows through (L in Figure 5.1) and the DC voltage applied to the inductor by the inverter. A larger inductance will yield a smaller $\frac{di}{dt}$ for a given voltage and so the slope of the sawtooth waveform in Figure 5.2 will be smaller. Brod and Novotny (Brod and Novotny, 1985) present an expression for the maximum switching frequency of a hysteresis controller, and this is given in Eqn. (5.1), where h is the hysteresis limit, and so the Hysteresis Band, ΔI, is equal to $2h$.

$$f_{sw(\max)} = \frac{V_{DC}}{9hL} \tag{5.1}$$

5.3 A Digital Hysteresis Current Controller

The active filter controller is implemented with a Digital Signal Processor (DSP) as presented in §4.2. All calculations are made digitally and the desired output current is represented numerically inside the DSP. The power switches in an inverter can either be on or off, and thus can be considered to be digital. The digital inverter controller interfaced the DSP

to the inverter with digital signals, rather than analogue voltages or currents. The Inverter Controller calculates the hysteresis limits mathematically and performs the magnitude comparisons with digital logic.

Traditional hysteresis current controllers have performed these functions with analogue operational amplifiers and comparators. Laying an analogue circuit out on a printed circuit board takes a considerable amount of time and uses a large number of components, many of which require fine adjustment. Any changes to the design would result in the procedure being repeated. A Field Programmable Gate Array (FPGA) is reprogrammable and allows modifications to the inverter controller to be made internally without any changes to the printed circuit board. An FPGA implementation is very compact because it is a single component and does not require a large number of support ICs (Retif *et al.*, 1993).

5.3.1 Previous Work

Other digital inverter controllers have been built with a DSP and programmable logic, usually erasable programmable logic devices (EPLD). One such example is that of Lee *et al.* (Lee, Suh and Hyun, 1996) where space voltage vector control was implemented with a Texas Instruments TMS320C31 DSP and an EPLD. Other researchers have implemented a digital hysteresis current controller using the TMS320C31 DSP (Li, Jin and Joos, 1995). The DSP's main task was to determine the harmonic content of the load current, with its spare processing time used to control the switches in a power inverter. There was sufficient surplus processing time to sample the injected current at 4800Hz, giving 81 decisions per cycle of 60Hz alternating current. A publication by Rahman *et al.* (Rahman *et al.*, 1997) claims that a sampling frequency of 1kHz is adequate for Digital Hysteresis Current Control. No justification for this claim was provided and the results of §5.3.3 show that 1kHz is almost certainly not a fast enough sampling frequency.

Recently published work (Tzou and Hsu, 1997) gives details of an FPGA implementation of a Space Vector Controller. This PWM technique is different to Hysteresis Current Control, however the FPGA implementation using a Xilinx XC4010 also gives very good high speed performance.

5.3.2 Advantages of Digital Control

All communications between the DSP and the FPGA are at 5V logic levels, as are the switching signals from the FPGA to the inverter. Having all external communication operating at 5V, rather than small analogue signal voltages, increases the immunity to

interference. Figure 5.3 compares the relative magnitudes of signal and noise for large and small signals, in an analogue and digital representation. The ‘Large Current’ signal has a much larger magnitude than the noise level and consequently there is little interference. However, the ‘Small Current’ is smaller in magnitude than the noise, so the interference will be high. When these currents are represented digitally (with the sample taken at the circle) the magnitude is greater than the noise for both large and small currents. This ensures that the susceptibility to noise is the same for all currents when represented digitally. Any feedback currents are analogue quantities and must be converted to a digital representation by an Analogue to Digital Converter (ADC) at some stage. It is most important that there is as little noise as possible at the input to the ADC as any errors introduced at this point will remain in the system.

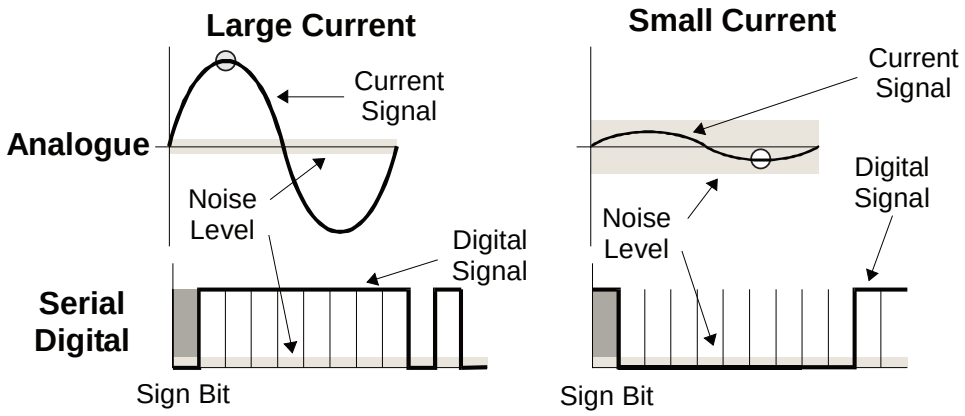

Figure 5.3 Comparison of analogue current signalling to digital signalling in the presence of noise.

5.3.3 The Effect of Sampling Frequency

Hysteresis current control requires current feedback. The sensed current is compared to the hysteresis limits generated from the reference signal and the desired hysteresis band. In an analogue system the comparisons are made continuously and the current will be forced to stay within the hysteresis band at all times while the inverter has sufficient power to control the current. With a digital controller, sampling occurs at discrete intervals. The sensed current is digitised and all comparisons are made digitally. The current information is updated at the sampling frequency of the ADC that samples the current feedback. If this sampling frequency is too low there is a chance that the current will have exceeded the hysteresis limits by the

time the next comparison is made. Figure 5.4 shows the results of a simulated inverter controller with three different sampling frequencies. The injected current is sampled at 260kHz in (a), at 26kHz in (b) and at 5kHz in (c). The hysteresis band, ΔI, has been set to 0.8A in all three cases. It can be seen that the current regularly exceeds the hysteresis band when the current feedback is sampled too slowly. The reference and the two hysteresis limits are shown as discrete samples in Figure 5.4. Once a switching decision has been made the transistor switches are left unchanged until the next sample is taken. This new sample is used by the decision logic to determine the next switching state. The centre solid line in each figure is the current reference and the jagged solid line is the inverter's current output. Two dashed lines represent the hysteresis limits used for the simulation.

Once the injected current travels outside of the hysteresis limits it will continue to do so until the next sample is taken. It is only at this point that the transistors are set so the current is ramped in the opposite direction. The high speed sampling of 260kHz gives very good results, and with only 3.8µs between samples the current will not get far outside the hysteresis limits. Compare this with the case of 5kHz sampling, where the current is regularly outside of the hysteresis limits. The time between samples has extended to 200µs, leading to extremely large current excursions. The TMS320C31 based controller (Li *et al.*, 1995) operated with a sampling rate of 4.9kHz, and this would give rather large current overshoots. The minimum sampling frequency needed to achieve an acceptable overshoot is a function of the injection inductance used and the voltage applied across the inductance. Each of these affects $\frac{di}{dt}$, and therefore the level of current overshoot. A suitable injection inductance was found to be in the range from 500µH to 10mH.

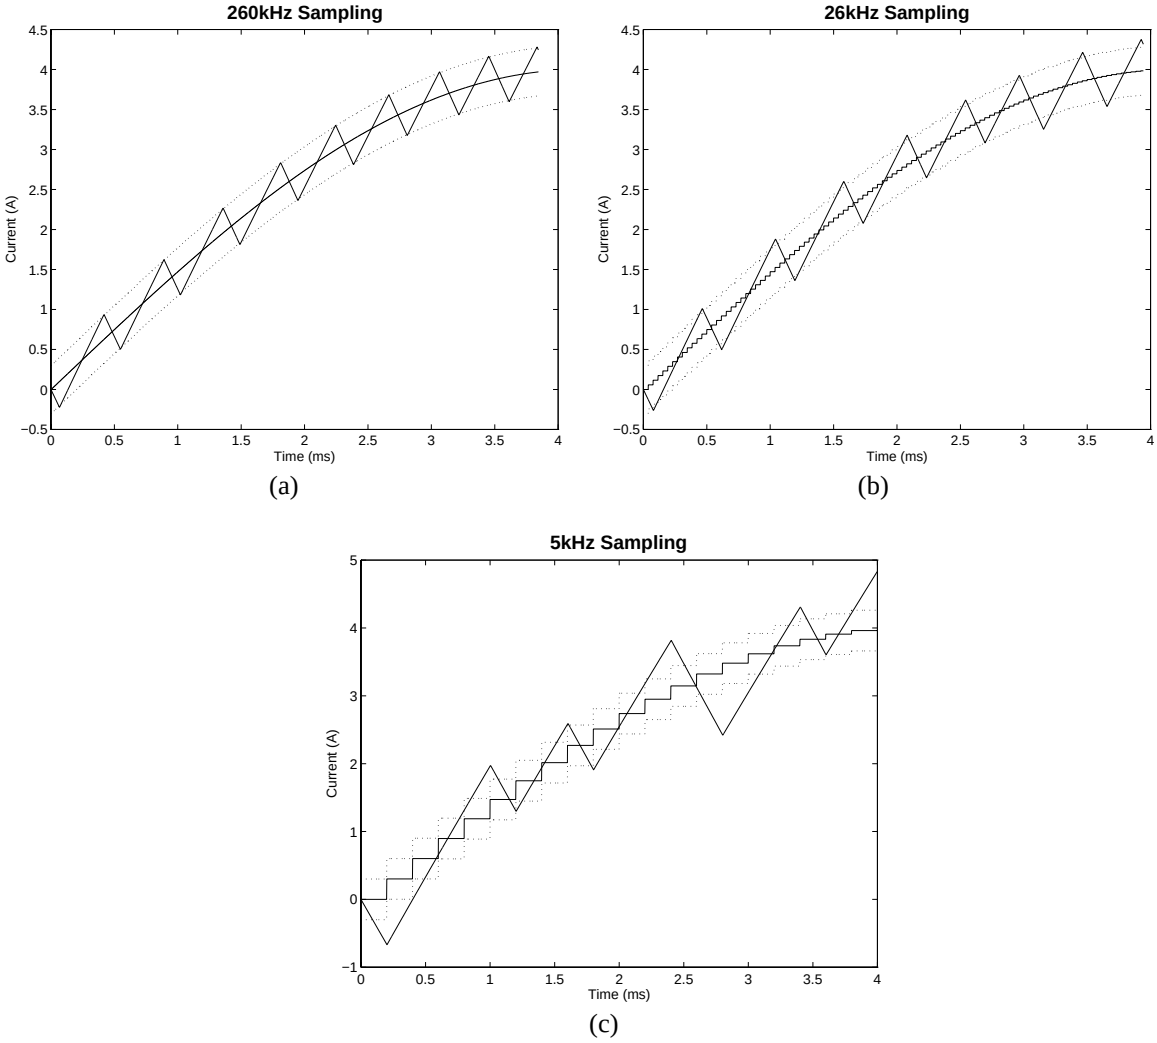

Figure 5.4 The effect of different sampling frequencies for the inverter feedback.

The smaller the injection inductor, the greater the $\frac{di}{dt}$, and this will increase the current overshoot. The worst case is where the current is just inside the hysteresis band when the comparison is made. The current will then continue on past the limit and will only reverse direction at the next sampling point. Eqn. (5.2) gives the maximum current overshoot Δi_{Over} for an inverter with a DC voltage V_{DC} injecting into an inductance L with a sampling frequency of f_{Sample}, assuming that $v_{Supply} = 0$. Ideally there will be no current overshoot and the current ripple, ΔI, will be solely due to the hysteresis limits. This is not practical and there will always be a degree of overshoot. The sampling frequency should be such that the overshoot is 5-10% of ΔI, which was found to give acceptable current waveforms. This keeps the current ripple reasonably close to the design value for a traditional Hysteresis Current Controller.

$$\Delta i_{Over} = \frac{\frac{1}{2}V_{DC}}{L}\frac{1}{f_{Sample}} \tag{5.2}$$

5.4 Controller Simulation

The operation of the proposed digital hysteresis current controller has been modelled using Simulink and MATLAB. A single phase model was developed since the three phases of the inverter must be able to operate independently. A neutral is provided and this prevents the situation arising where the hysteresis band can double due to three phase interactions (Brod and Novotny, 1985).

The block diagram of the system using the hysteresis controller (the reference generator, the controller, the inverter and the inductor) is shown in Figure 5.5. To get the best possible simulation of a real digital hysteresis current controller the currents were quantised and sampled. This gives the simulated controller the same dynamic characteristics as a controller implemented in an FPGA. The hysteresis controller makes switching decisions based on the reference current signal and the current feedback. The switching signals are used to control a model of a three-phase inverter that in turn produces a three phase voltage output. The voltage is applied across the injection inductor model, which through the use of differential equations determines the current that flows through the inductor. This injection current is the output from the simulation and is recorded for further analysis.

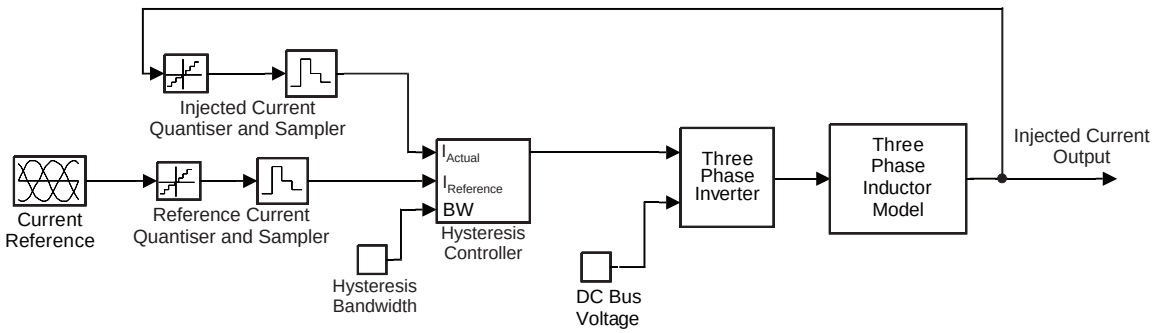

Figure 5.5 Simulink system for Hysteresis Controller simulation.

Further explanation of the operation of the Hysteresis Controller is given by Figure 5.6. The Hysteresis Controller takes three inputs: the desired reference current signal, the

Hysteresis Bandwidth and the current injected by the inverter. The controller is closed loop because the 'Injected Current Output' is dependant on the operation of the inverter controller. The Hysteresis Band Generator takes the reference signal and the desired bandwidth and creates the Lower Hysteresis Limit and the Upper Hysteresis Limit (as shown in Figure 5.2). Equations (5.3) and (5.4) give the expressions for the upper and lower limits.

$$I_{Lower} = I_{reference} - \tfrac{1}{2} I_{bandwidth} \tag{5.3}$$

$$I_{Upper} = I_{reference} + \tfrac{1}{2} I_{bandwidth} \tag{5.4}$$

Magnitude comparators determine when the 'Actual Current' exceeds the upper limit or is less than the lower limit. If the injected current does exceed the upper hysteresis limit then a negative voltage has to be applied to that phase to reduce the current. This is achieved by turning on the bottom switch in the appropriate leg of the inverter. The SR latches hold the state for each inverter leg. A "1" means that the upper switch is on and a "0" means the lower switch is turned on. The "S" input sets the output to "1" and the "R" input sets the output to "0". A 2µs propagation delay is used to model the dead-time protection used in a real inverter controller. The three outputs are the driving signals for the top switch in each leg of the inverter. The signal for the bottom switch in each leg is the logical NOT of the top switch signal.

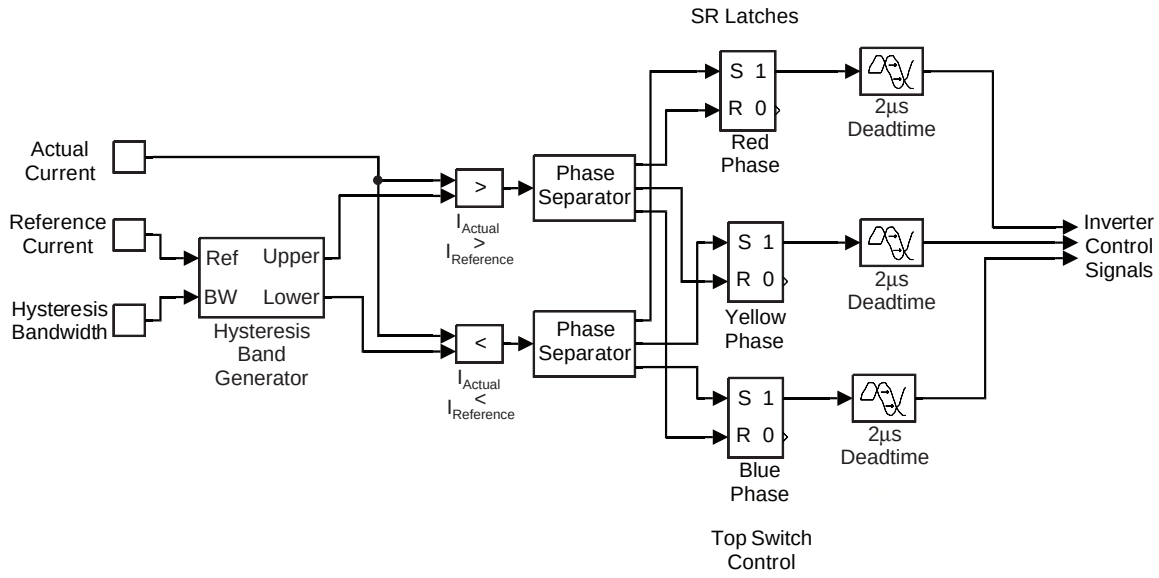

Figure 5.6 Simulink implementation of a Hysteresis Current Controller.

5.4.1 Sinusoidal Performance

The current overshoot examples in Figure 5.4 were generated with the simulated hysteresis current controller. The performance of this controller is shown in Figure 5.7 where a sinusoidal reference is supplied to the hysteresis current controller. The reference signal frequency is 36Hz, the DC supply is 60V and the inductance is 9mH. A 36Hz reference was used to show that the inverter is capable of injecting currents at frequencies other than 50Hz. The peak current is 6A and the hysteresis band is 0.2A. It can be seen that the switchings give a current that follows the sinusoidal reference. Figure 5.7(a) is scaled to show the full range of the sinewave. The peak to peak value of this current waveform is 12A, and so the current ripple is only 1.7%. The frequency spectrum of the current in Figure 5.7(b) shows that the majority of switching noise is in the band 5–7kHz. The maximum expected switching frequency, as given by Eqn. (5.1), is 7.4kHz which agrees with the spectrum in Figure 5.7(b). Figure 5.7(c) illustrates the current ramping up and down in more detail by looking at a 1.6A range rather than a 12A range. The slope of the current ripple for this waveforms is 6.7A/ms. The width of the Hysteresis Band is quite clearly shown at the time of 6ms in Figure 5.7(c). The current ranges from 5.8A to just over 6.0A and is very close to the 0.2A design value.

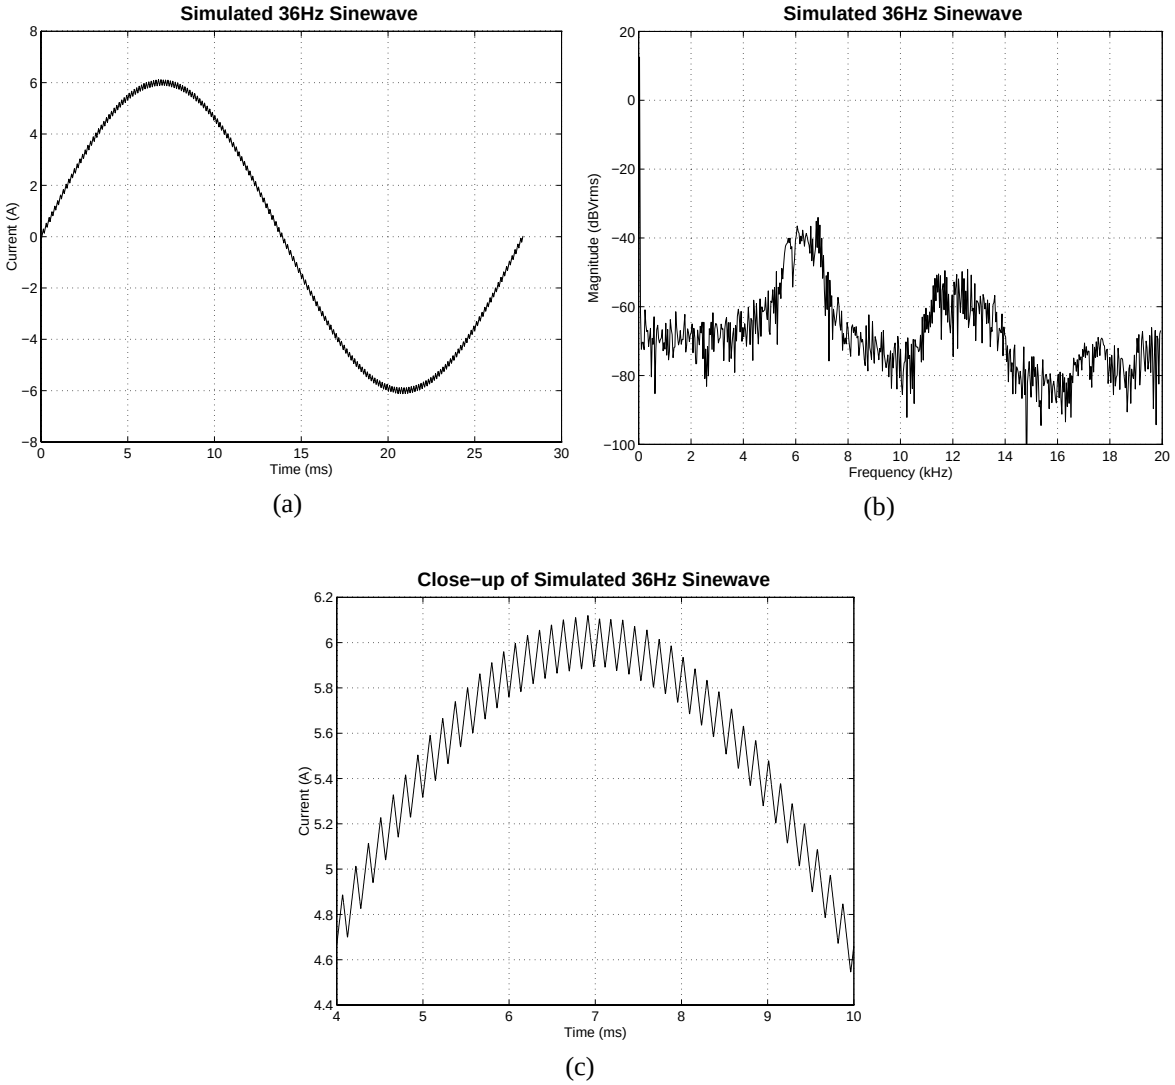

Figure 5.7 Simulation of a 36Hz sinewave from Simulink in (a) the time domain and (b) the frequency domain. Detail of the current waveform is shown in (c).

5.4.2 Harmonic Current Output

This digital inverter controller is intended for use in an active filter, and so a sample compensating current, shown in Figure 5.8(a), was generated. The fundamental frequency is 50Hz and the injection inductance is 4mH. The compensating current is scaled for an overall load current of $0.71A_{RMS}$ and the hysteresis band is 0.2A. The slope of the current is much higher for this waveform than for the pure sinusoid because of the smaller injection inductance.

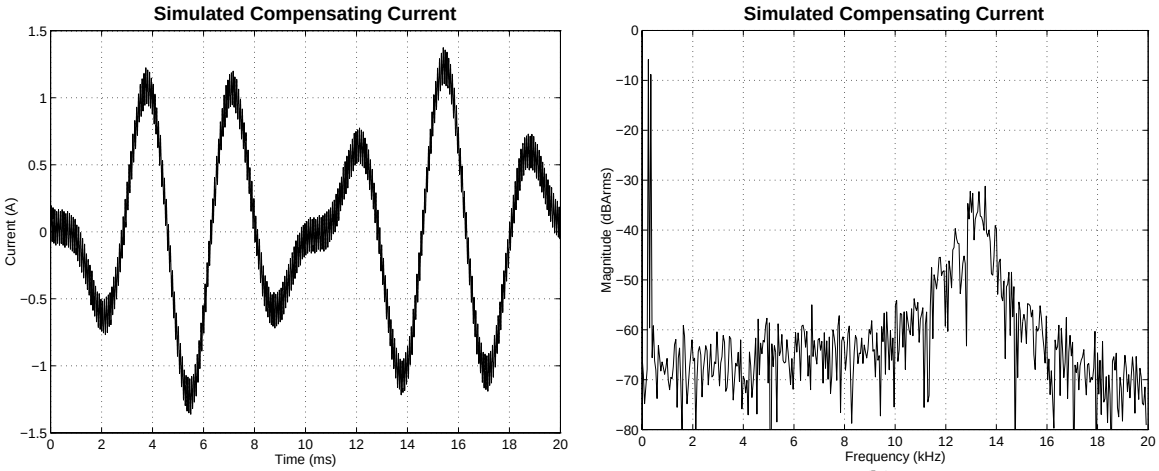

(a) (b)

Figure 5.8 Simulated compensating current output from Simulink.

Figure 5.8(b) shows the frequency spectrum of the current in Figure 5.8(a). The two prominent frequency components at 250Hz and 350Hz are the harmonics present in the reference signal. The concentration of signal at 13.5kHz is the high frequency switching signal from the inverter. The maximum expected frequency in the inverter output is 16.7kHz and this agrees with the frequency spectrum in Figure 5.8(b). The unwanted frequency components can be filtered by an LC low pass filter because they are so far removed from the frequency band of interest, which is limited to 2.5kHz.

If an injection transformer is used, a capacitor is installed on the high voltage side of the transformer. However, if an injection inductor is used, a capacitor must be placed on the opposite side of the inductor to the inverter. These connections are illustrated in Figure 5.9 for inverters using (a) injection transformers to get the high voltage output and (b) inverters that inject directly into the power system through an injection inductor.

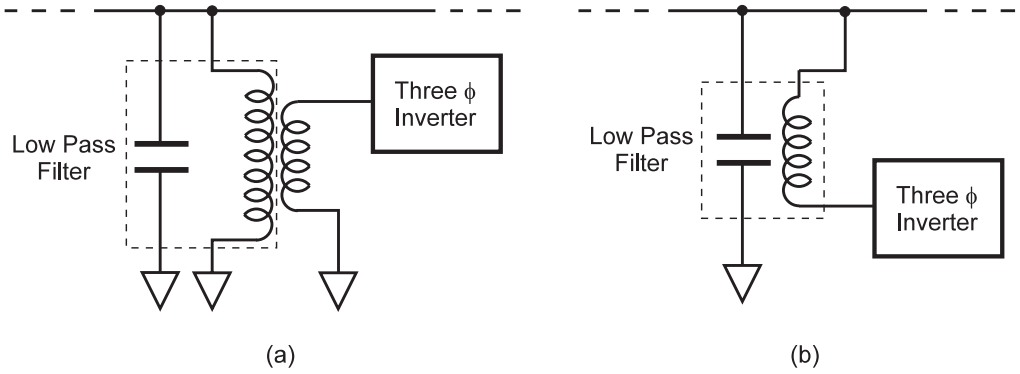

Figure 5.9 Placement of the filter capacitor used to remove the switching noise from the inverter output current.

5.5 Hardware Implementation in an FPGA

The digital inverter controller interfaces the digital signal processor to the three phase inverter. Figure 5.10 shows the inverter system with the controller, the power inverter and current feedback. A LEM current sensor and transimpedance amplifier convert the high current output of the inverter to a voltage suitable for digitising by the ADC. The reference current from the DSP and the digitised injection current are transmitted digitally to the FPGA. Switching decisions are made by the FPGA that control the output voltage of the inverter. The voltage applied to the inductor causes a current to flow into or out of the power system.

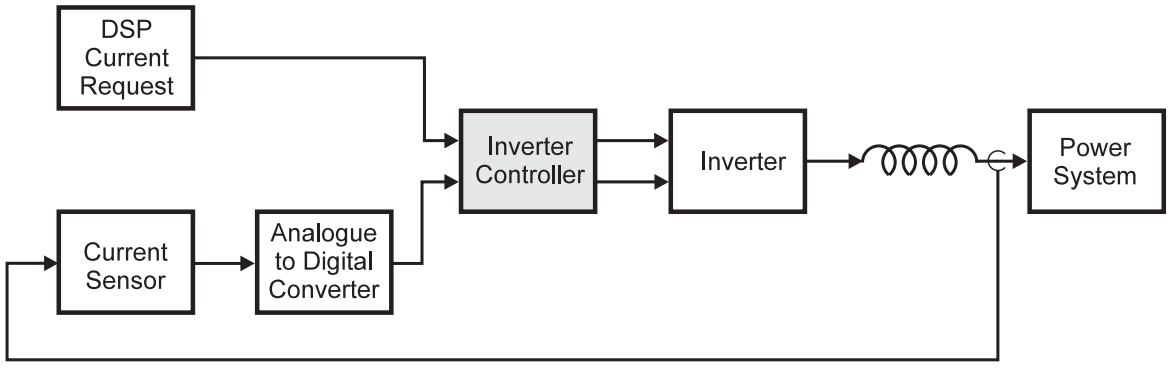

Figure 5.10 System diagram of the inverter and controller for one of three phases.

Figure 5.11 expands upon the inverter controller, highlighting the internal blocks used for hysteresis current control. The controller is completely digital and is implemented with the FPGA.

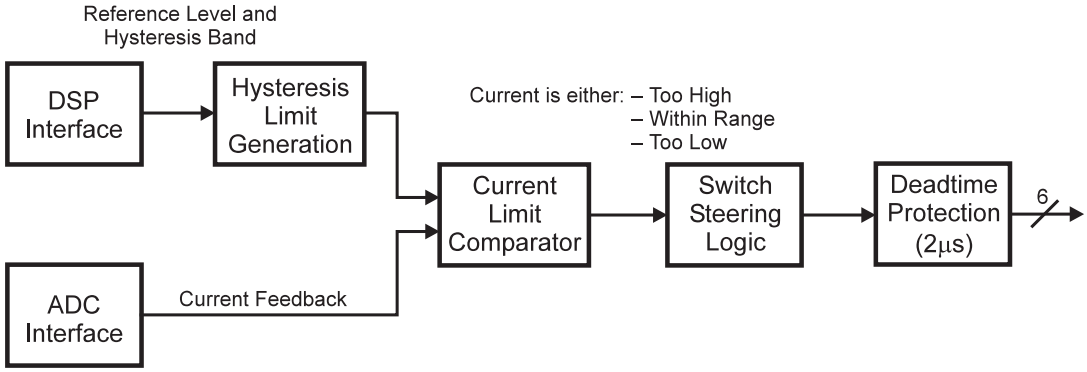

Figure 5.11 Block diagram of the inverter controller.

The DSP communicates with the FPGA through its synchronous serial port with an 8MHz clock. A Maxim MAX176 12 bit serial ADC is used for the current feedback with a sampling frequency of 260kHz. Serial communication reduces the number of input connections to the FPGA but increases the internal logic requirements. The lower number of connections makes opto-isolation easier should non-isolated current transducers, such as resistive shunts, be used.

Two adder/subtracters per phase generate the hysteresis limits from a reference and hysteresis band supplied by the DSP. Signed magnitude comparators check the current feedback against these limits. The 'Switch Steering Logic' block uses sequential logic to determine the appropriate switching pattern so the current is driven in the appropriate direction through the injection inductor. Deadtime protection is performed by a state machine that ensures a 1.6μs delay between switchings. This allows for the long current tail of an IGBT and prevents conduction faults that would destroy the inverter. A Xilinx XC4010E FPGA was used to implement the phase controller. This FPGA is approximately equivalent to 10000 gates (Xilinx, 1996) and 37% of its resources were used.

5.6 Experimental Results

The experimental implementation of a digital hysteresis current controller was tested with sinusoidal and harmonic current outputs. Where possible, the experimental outputs were matched to those used in the Simulink model of the current controller.

5.6.1 Sinusoidal Outputs

The current waveform is the same as the simulated waveform from §5.4.1. The fundamental frequency is 36Hz and the magnitude is 6A, with a hysteresis band of 0.1A. The

DC power supply for the inverter is $60V_{DC}$ and a 9mH injection inductor was used. Figure 5.12 shows the spectrum of the injected current and the time domain waveform.

The upper plot of Figure 5.12(a) shows the frequency spectrum of the current output. There is a large peak near 0Hz representing the fundamental frequency. The asynchronous nature of the hysteresis current controller is shown by the band of frequencies at 4.7kHz.

The lower plot in Figure 5.12(b) shows just over one cycle of the output current shown in the time domain. The current ripple is small in comparison to the overall magnitude of the current and that there are no significant current overshoots. Overshoots have occurred and Figure 5.12(b) shows one of these in one of the positive half cycles of current. The circle -1.2ms from the centre shows where the current has deviated from the reference by twice the design limit for the hysteresis controller. This was due to the inverter controller sampling the injection current when it was just inside the hysteresis limit. The hysteresis band in Figure 5.12(b) is relatively constant over the half cycle. This indicates that the sampling frequency is high enough as there are no large current overshoots.

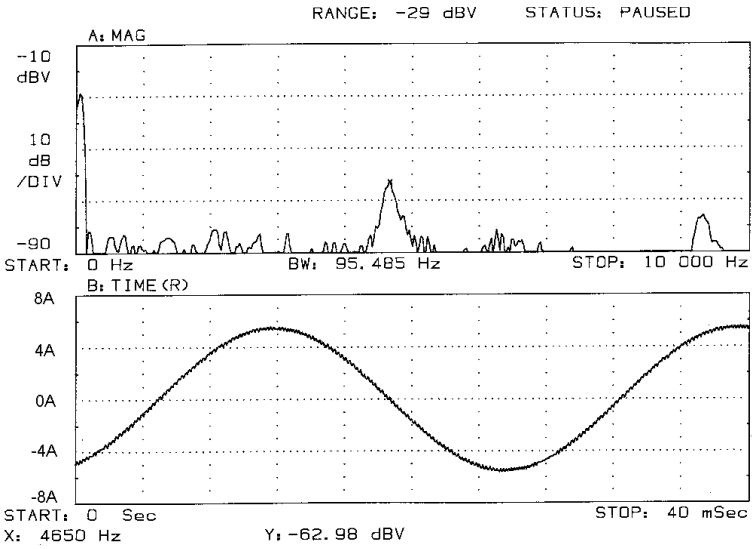

(a)

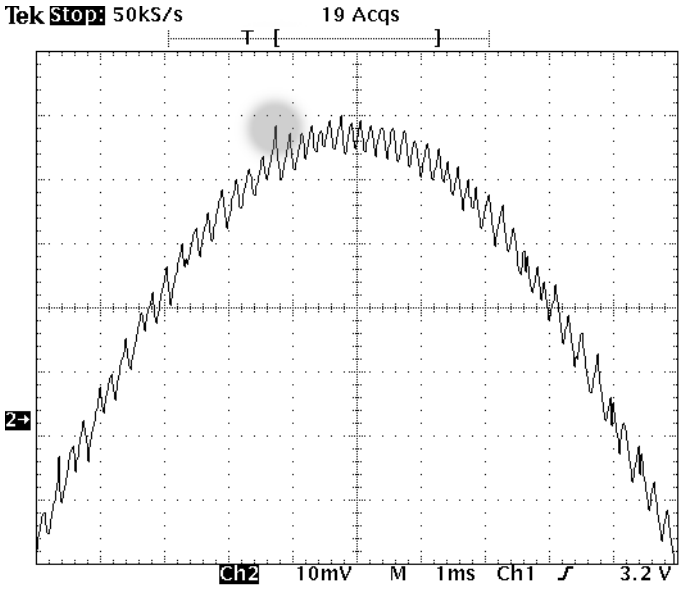

(b)

Figure 5.12 Sinusoidal current outputs from the digitally controller inverter of (a) the time and frequency domain Plot of the inverter output current, and (b) close up detail of the hysteresis in the current output from the inverter

5.6.2 Harmonic Current Outputs

The compensating current that would result from a three phase bridge rectifier with capacitive smoothing was generated artificially by the DSP. This is the same test waveform as was used for the Simulink simulations in §5.4.2. The fundamental frequency is 50Hz and a 4mH injection inductance was used. The overall output current is $0.71A_{RMS}$, with a hysteresis band of 0.2A used for this test.

The upper plot in Figure 5.13 shows frequency spectrum of the ‘compensating current’. The wide peak at the left of the spectrum contains the 250Hz and 350Hz components of the compensating current. The other frequency components are at least 10dB down from the desired output.

The lower waveform illustrated in Figure 5.13 is the time domain representation of the output current. The $\frac{di}{dt}$ of the output current has been increased by reducing the injection inductance so the current can follow the higher frequency reference. This gives more switchings at the top of each current peak when the slope is near zero. This is seen in Figure 5.13 as increased current noise at the top of the peaks. The switching noise cannot be seen in the frequency spectrum because the maximum displayed frequency is 10kHz.

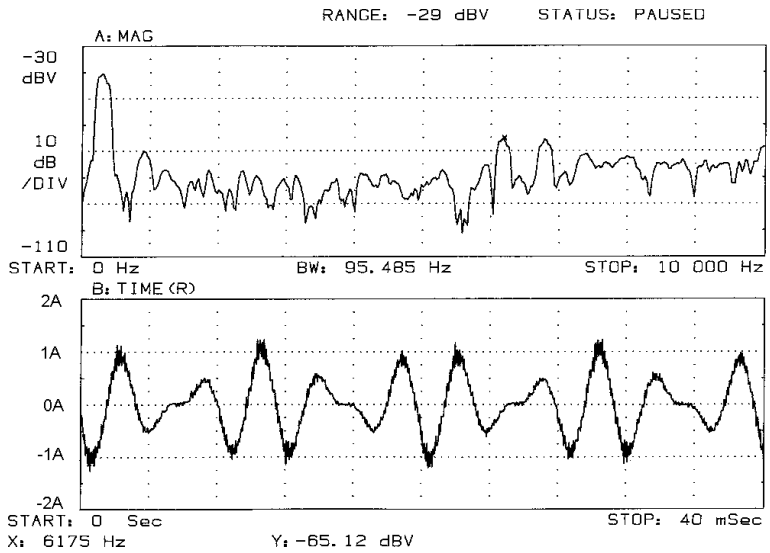

Figure 5.13 Compensating current output from the inverter.

5.7 Discussion

5.7.1 Average Switching Frequency

The hysteresis current controller ideally has a constant current ripple, ΔI, but no defined switching period, Δt (Brod and Novotny, 1985). Further tests were conducted with a single 60Hz sinusoid injected into the inductor. Table 5.1 shows the frequency and average magnitude of the switching component of the current with respect to the magnitude of the

60Hz fundamental. The entries where the hysteresis limit is 10% of the peak current are highlighted. The inverter bus voltage was a constant $60V_{DC}$.

Table 5.1 Relative magnitude and frequency of switching component.

I_P	h (A)					
(A)	0.1	0.2	0.3	0.4	0.5	0.6
1.0	7.2% 7.1kHz	15.7% 3.7kHz	21.4% 2.5kHz			
2.0	2.5% 7.7kHz	6.1% 3.5kHz	10.2% 2.5kHz			
3.0	1.8% 7.8kHz	2.2% 3.7kHz	3.6% 2.4kHz			
4.0	0.8% 7.7kHz	1.8% 3.7kHz	2.3% 2.4kHz	3.3% 1.8kHz		
5.0	0.7% 7.7kHz	1.3% 3.6kHz	1.8% 2.4kHz	2.3% 1.8kHz	3.7% 1.4kHz	
6.0	0.5% 7.7kHz	0.6% 3.6kHz	1.0% 2.4kHz	1.8% 1.7kHz	2.7% 1.3kHz	3.1% 1.0kHz

The results in Table 5.1 show that the magnitude of the switching component is, on average, 4% of the fundamental current when the hysteresis limit is one tenth of the reference current peak. The lowest level of the switching component was with h = 0.1A, I_P = 6.0A and was 0.5% of the fundamental and is an improvement over the 3% magnitude obtained with h = 0.6A. The average switching frequency is shown by the results in Table 5.1 to depend only upon h, as is expected from Eqn. (5.1).

5.7.2 Noise Immunity

One benefit of digital control is that the current reference is transmitted serially from the DSP to the inverter controller. It was predicted that this would enhance the controller's ability to withstand electromagnetic interference (EMI). The susceptibility to EMI of this prototype inverter controller was tested by operating a 146.475MHz radio transmitter at a variety of distances from the controller. Even though oscilloscope waveforms showed

increased noise on the voltage signal feeding to and from the FPGA, there was minimal disturbance to the current waveform from the inverter.

Disturbance only occurred when the tip of the transmitting antenna was touching the Xilinx FPGA. The inverter failed to operate once the transmitter was stopped, and only started again when the DSP updated the reference and hysteresis limits. This suggests that a register change occurred in the Xilinx due to the high electrical fields present (about 100-200V/m). This level of EMI would seldom be found except near radio transmitters. It can be assumed that the digital inverter controller itself is not susceptible to EMI.

The one weakness in the system is the analogue to digital converter (ADC) used for digitising the current feedback. The input to the ADC does show noise in the presence of EMI and so great care must be taken to ensure that this device and associated analogue circuitry are well shielded. Figure 5.14 shows the voltage present at the input to the MAX176 analogue to digital converter. High frequency noise is present at a level of approximately 40mV, which equates to four bits of resolution of the ADC. Switching transients also appear in the waveform when the current changes direction, but with improved circuit layout these can be minimised.

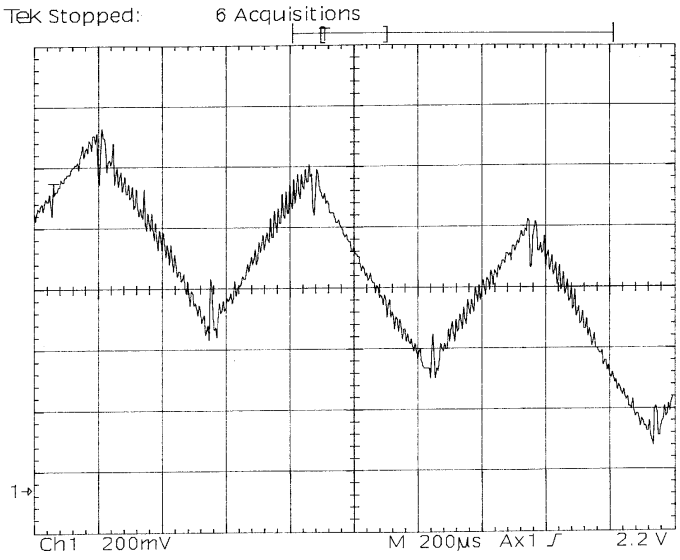

Figure 5.14 Voltage waveform at the current feedback ADC input.

5.7.3 Current Overshoot

The sampling frequency for the current feedback was fixed at 260kHz in the real system, and therefore some current overshoot did occur. The current swing was measured while the switching frequency tests were being performed. For this test, L = 9mH, V_{DC} = 60V and f_{Sample} = 260kHz. From Eqn. (5.2) the value of overshoot, Δi_{Over}, is 13mA. Table 5.2 summarises the predicted level of maximum overshoot for six values of h.

Table 5.2 Maximum current overshoot out of the hysteresis band.

h (A)	**Overshoot** (mA)	
	Predicted	Actual
0.1	13	28
0.2	13	40
0.3	13	74
0.4	13	76
0.5	13	65
0.6	13	60

Whilst the predicted maximum overshoot is independent of the hysteresis limit, h, this is not so for the experimental results. Measuring small currents is difficult and the input to the current feedback ADC (shown in Figure 5.14) has a noticeable level of noise. This will affect the inverter controller's ability to determine accurately when the current has exceeded the hysteresis band. It is expected that by dealing with larger currents the effect of noise on the analogue circuitry will be reduced.

5.8 Summary

A novel digital hysteresis current controller has been designed and implemented in a Xilinx Field Programmable Gate Array. This completely digital controller makes the switching decisions for hysteresis control and provides deadtime protection. The theoretical performance of a digital hysteresis controller has been presented. The inverter controller has

been tested on a small scale inverter, with experimental results providing performance data and verification of the accuracy of the simulation model used.

The digital implementation of a hysteresis current controller was successful and performed as expected. High levels of immunity to EMI make this type of current controller very promising for use in full size high frequency switching inverters. A digital interface makes connection to an active filter controller straightforward and robust. The controller is adaptable for use in any application requiring current control of a voltage source inverter and is not limited to active filtering.

The next chapter describes the software that ran on the hardware implementation of the active filter controller. The different harmonic isolation methods are discussed, along with the general purpose ‘framework’ that was common to each.

Chapter 6
Software Development

6.1 Software System

Chapter 4 discussed the hardware developed and used for this implementation of a digital active filter controller. The DSP's software allowed this active filter controller to operate with more flexibility than a traditional analogue controller. While this software can operate in isolation, a personal computer (PC) was used to display the measurements made by the DSP. This chapter describes the function and structure of the DSP software, the PC control software and the communications between the two.

6.2 DSP and PC Software Interface

The Sonitech Spirit-30 'C30 development board (described in §4.2.1) has 32k × 32 bits of dual port static RAM. Dual port RAM chips have two sets of busses and control signals. This allows two devices to access the same memory simultaneously and is often used when two microprocessors need to access the same memory without interfering with each other.

The DSP program was executed from the dual port RAM. The ISA bus interface on the Spirit-30 board allowed the PC to read from, and write to, the dual port RAM with a software library provided by Sonitech. This library was compiled into the main PC software and these routines were used to read from and write to variables in the DSP's memory while the DSP was executing its program. Loading the DSP software into the dual port RAM meant that changes to the code could be made extremely often and easily, removing the need to extract, erase, program and insert EPROM chips into the development board.

6.3 Digital Signal Processor Software

The DSP software was written in a combination of C and assembly code. C is particularly suited to the TMS320C30 because of the DSP's large register set and software stack. The Texas Instruments C compiler takes advantage of the DSP's special features, such as parallel instructions and delayed branches (Simar and Davis, 1988). C is rapidly becoming the preferred language for embedded software development (Chassaing, 1992). Programming

in a high level language such as C, instead of assembly language, tends to increase the size of the executable program and reduces the execution speed. Assembly language is generally faster but much harder to document and maintain (Chassaing, 1992). The approach adopted for this software was to write the majority of the program in C, with the time critical routines written in assembly language modules that were called from the main program. The digital filter routines discussed in §7.3 are examples of these assembly language modules. The DSP software consisted of a main program that executed continuously and a set of interrupt handlers to deal with internal and external interrupts.

6.3.1 Main Program

The main DSP program is responsible for initialising the DSP and its on chip peripherals. Buffers and other variables that are used by the active filter controller must also be initialised before active filtering starts. A flowchart showing the events that occur when the DSP starts executing code is shown in Figure 6.1. The DSP first enabled the 64 word program cache that increases the processing speed of the DSP. The Primary Bus defaults to having seven ‘waitstates’ (clock cycles where the DSP waits for slow external devices) to allow for the slowest possible external memory. The dual port RAM is high speed memory and does not require any waitstates. Version information (version number, time and date of compilation) is embedded into the DSP software to be retrieved and displayed by the PC software. This ensured that the correct version of the DSP software was running and showed the harmonic isolation technique being evaluated. Other variables such as storage buffers and data pointers were initialised to known states in readiness for the start of normal program operation. The two synchronous serial ports (SP0 for the Inverter and SP1 for the MAX186 ADC) were then configured and the three interrupts used (INT1, INT2 and RINT1) were enabled in the Interrupt Enable (IE) register. Once configured, the DSP paused until the PC software indicated that it was running and ready to take results from the DSP.

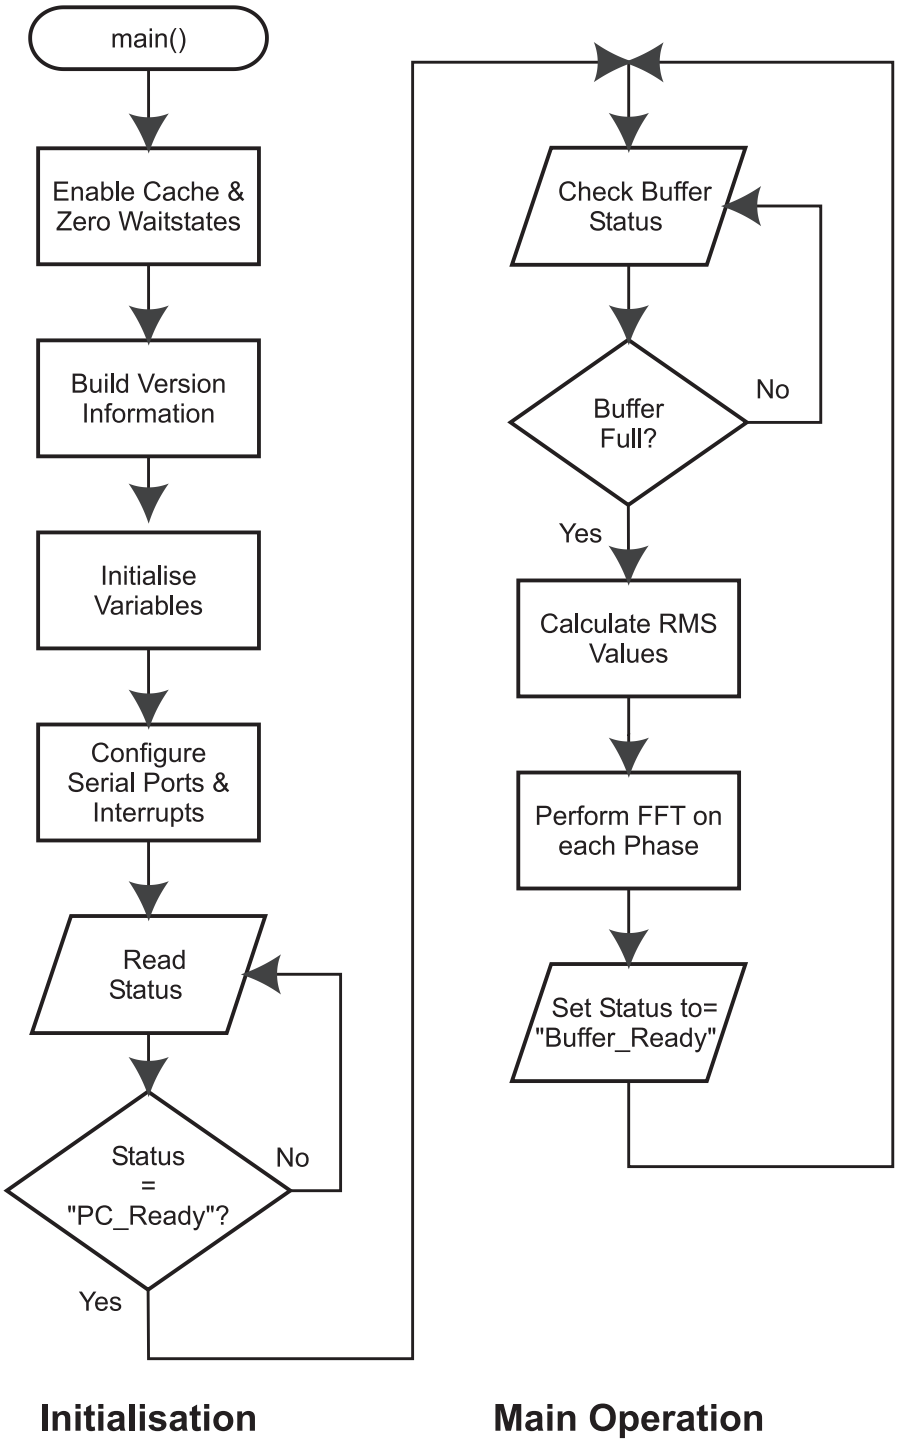

Figure 6.1 Flow diagram of the main DSP program.

The status of the input buffer was regularly polled to see when 128 samples (a complete mains cycle) had been accumulated. Once a complete data buffer was acquired, the following power quality measurements were made on each of the three phases: RMS current, Total Harmonic Distortion (THD) and the frequency spectrum. Data transfer from the DSP to the PC was through a dedicated memory buffer that contained raw integer samples for each

phase, along with the RMS and THD measurements. Data transfer was initiated by the PC control software once the DSP indicated measurements were ready.

Figure 6.2 gives a graphical breakdown of the power quality calculations by showing the calculation stages and the relative time taken for each. The frequency spectrum of each phase is calculated first, taking approximately 3.9ms for all phases to be calculated. Once this was achieved, the RMS value for each phase was determined, taking a total of 0.3ms.

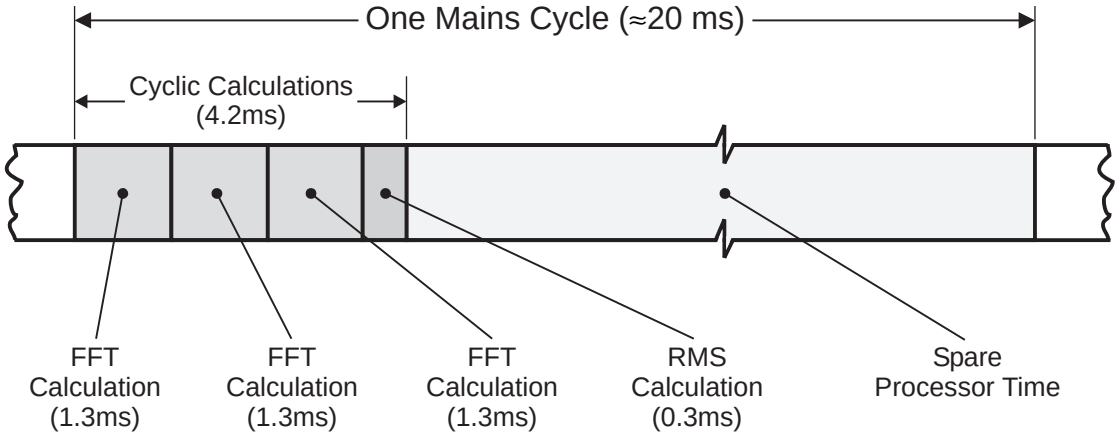

Figure 6.2 Timing diagram of the DSP execution for one mains cycle

Figure 6.3 gives more detail on the steps taken to calculate the frequency spectrum of each phase. The first step was to make a copy of the integer samples from the ADC for display on the PC screen as a waveform. The integer samples were then converted to a floating point format for further calculations. This conversion scaled the integer values with a range from -2048 to +2047 to the true current range of $-35A_{Peak}$ to $+35A_{Peak}$. The assembly language Fast Fourier Transform routine (Tessarolo, 1991) was called by the C program to calculate the complex frequency spectrum. The magnitude of each frequency component was obtained by taking the square root of the sum of the real and imaginary components squared, as given by Eqn. (6.1). These frequency magnitudes were used to calculate the THD with the equation given by Eqn. (2.3) and were converted to integer values for display by the PC software.

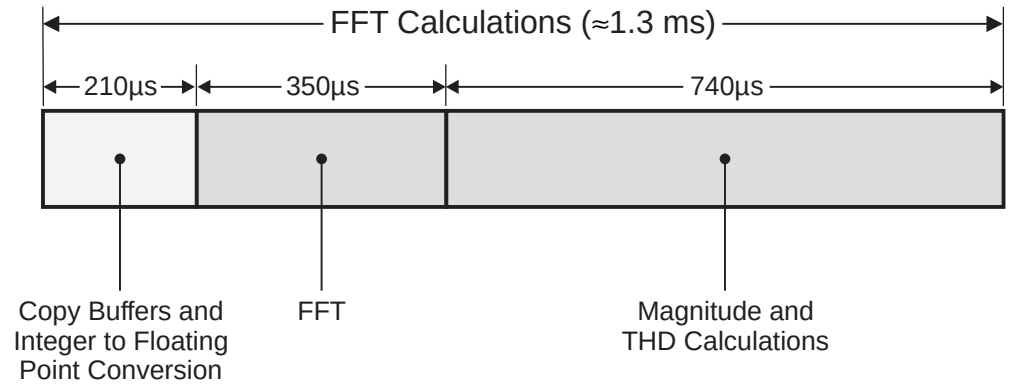

Figure 6.3 Timing diagram with expanded detail on the calculation of the frequency spectrum.

$$|F_i| = \sqrt{\mathrm{Re}\{F_i\}^2 + \mathrm{Im}\{F_i\}^2} \tag{6.1}$$

While the above calculations were being performed, interrupt service routines regularly broke into the program, performed specific tasks, and then returned to the original point. Further information on these routines is presented in the next section.

6.3.2 Interrupt Service Routines

Interrupts caused the DSP to jump out of the normal program flow and execute an Interrupt Service Routine (ISR) in response to some external or internal stimuli. In this active filter controller the stimuli were the Sample Interrupt, Zero Crossing Detect Interrupt and the Serial Port 1 Received Data Interrupt. Once the ISR was finished, the DSP continued execution from where it was before the interrupt occurred.

6.3.2.1 Zero Crossing Detect (INT2)

A flowchart illustrating the Zero Crossing Interrupt Service Routine is shown in Figure 6.4. The Zero Crossing Interrupt is used to synchronise the DSP to an absolute position in the mains cycle. The reference voltage for this controller was taken from Yellow Phase and provided the 0° phase reference. The Sample Rate Multiplier is not perfect and external disturbances may have caused the DSP to slowly drift from the phase reference. The Zero Crossing Detector corrects this by ‘zeroing’ the phase reference every mains cycle. Data acquisition is ‘double buffered’ by the DSP to allow the post processing (RMS, FFT and THD calculations) to occur while data is being collected. The Zero Crossing ISR manages the buffer swapping and ensures that background calculations (discussed in §6.3.1) are not

performed using the data buffer that is currently being written to. Before leaving this ISR a software flag is set to tell the main program that a complete mains cycle's worth of data has been collected and is ready for processing.

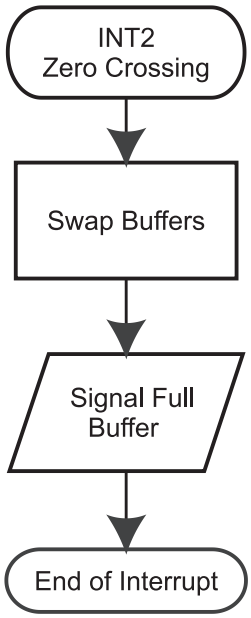

Figure 6.4 Flow diagram of the Zero Crossing Detect ISR.

6.3.2.2 Sampling Interrupt (INT1)

The Sampling Interrupt was used to make the DSP sample load currents at a regular interval, with the exact sampling frequency being 128 times the mains frequency. Interrupts were generated by the Sample Rate Multiplier (described in §4.3.3) which caused the DSP to execute the Sample Interrupt Service Routine. The flowchart for this routine is presented in Figure 6.5. When the Sample Interrupt occurred (at a nominal 6.4kHz) the DSP transmitted a request for the Analogue to Digital Converter (ADC) to sample Channel 1 (Red Phase), set a status variable to say that Channel 1 was requested and then exited the ISR. When the ADC conversion was complete, which took approximately 20µs, the result was shifted into Serial Port 1 on the DSP. Once sixteen bits were received the Serial Port generated an internal interrupt to tell the DSP that new data was waiting. The flag set by the Sample ISR told the Serial Port Received Data ISR where to store the new data after converting it from a signed 12 bit integer to a signed 32 bit integer. This latest reading was then used to calculate the compensating current (described in more detail in the next section) using the particular

technique under evaluation. The digital representation of the compensating current signal was then transmitted out of Serial Port 0 to the power inverter or simulator. If the current reading was not for Channel 3 (Blue Phase) the next channel was sampled by transmitting a sample request to the ADC. Once all three channels (Red Phase, Yellow Phase and Blue Phase) were sampled, the buffer position for the next set of samples was increased by one.

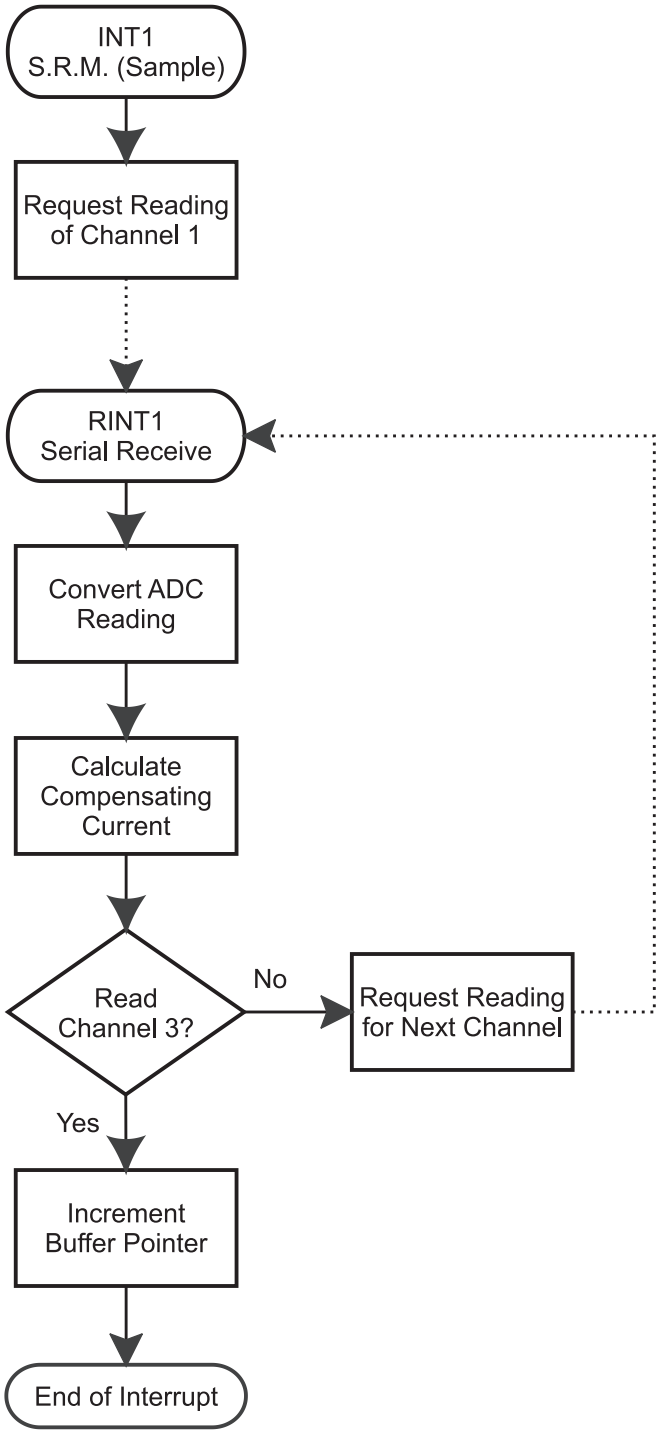

Figure 6.5 Flow diagram of the Sample Rate Multiply ISR.

The dotted lines in Figure 6.5 indicate that other processing takes place (the background task) before going to the block at the arrowhead. This ‘multitasking’ is shown diagrammatically in Figure 6.6, with the ISR functions shaded dark grey and the background processing shaded light grey. The diagram also shows the interrupt that caused each ISR to be called, either the Sample Interrupt (INT1) or Serial Port 1 Data Received (RINT1). The INT1 ISR is executes in approximately 3µs because it is simply transmitting a sample request to the ADC. The RINT1 ISRs for Red Phase and Yellow Phase were longer than the RINT1 ISR for Blue Phase (6µs instead of 4µs) because these transmitted a sample request for the next channel.

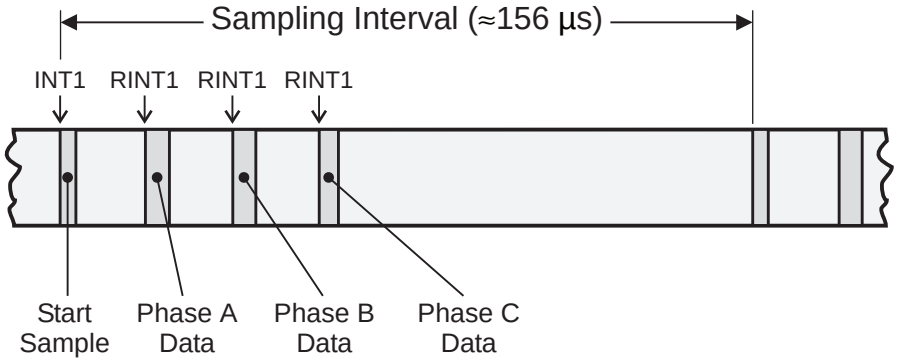

Figure 6.6 Timing diagram showing the sampling of each of the three phases at a nominal 6.4kHz.

6.3.3 Compensating Current Generation

Three types of harmonic isolation techniques were examined: digital filtering, sinusoidal subtraction and the Fast Fourier Transform. Each of these was performed differently and will be discussed separately.

6.3.3.1 Digital Filtering

The abbreviated code fragment that performed the isolation of harmonic currents and the generation of the compensating current is shown in Figure 6.7. The digital filtering routines were self contained, with the assembly code called with a new sample as the single parameter and the filtered value was the single return.

```
ti_float = (float)ti_value;              // Convert Int to Float
firfilt1(256, &ti_float, &filtfloat);    // Perform HPF on Phase 1
icomp = (int)(-filtfloat);               // Ic = -Ih
                 ⋮
shortint = (icomp + 2048) & 4095;        // Add a DC offset for DAC
while ((sp0[gcr] & 0x8) == 0x8);         // Wait until SP1 shifter empty
sp0[dtr] = shortint;                     // Transmit Ic to DAC
```

Figure 6.7 C code for the generation of compensating currents using digital filtering harmonic isolation techniques.

The C program listed in Figure 6.7 performs the following steps:

1. The 32 bit integer is converted to a floating point number in the range -2048 to +2047.
2. The filtering routine is called with the new sample as its argument. The design and implementation of these digital filters is described in more detail in Chapter 7.
3. The compensating current signal is the negative of the harmonic current signal produced by the digital filter routine.
4. A DC offset is added to the compensating current.
5. A check is made to ensure that the serial port transmitter shift register is not busy.
6. The digital representation of the compensating current is transmitted out of the serial port.

The steps described above are performed on every sample for each phase. A separate filter routine is required for each phase because the data is contained within the routine. The Notch Filtering harmonic isolation technique used an IIR digital filter to mimic the magnitude characteristic of an analogue Butterworth notch filter. A variation on the Notch Filtering technique was tried with two different high pass filters. These were implemented with linear phase FIR filters to give a net zero phase shift between the input and the output.

6.3.3.2 Sinusoidal Subtraction

The sinusoidal subtraction harmonic isolation technique is slightly more complex than the techniques that solely use digital filtering. The section of C code that performs the sinusoidal subtraction is shown in Figure 6.8.

```
firfilt( 256, &ti_float, &lpfcurrent ); // LPF (Get 50Hz Magnitude)

// Look at the peak of the sinewave for the value
if ( bpointer == 117 )                  // Positive half-cycle
     red_sine_mag = lpfcurrent;
if ( bpointer == 53 )                   // Negative half-cycle
     red_sine_mag = -lpfcurrent;

// Calculate the compensating current at this particular point
ih = ti_float - red_sine[bpointer] * red_sine_mag;
icomp = (int)(-ih);
```

Figure 6.8 C code for compensating current generation using Sinusoidal Subtraction

The load current is low pass filtered to obtain the 50Hz component of the current. The group delay of the filter is exactly one cycle because a 256 tap FIR filter is used (Chassaing, 1992). The peak of the 50Hz current waveform will always be at the same point in the buffer due to the Zero Crossing Detect synchronisation. Table 6.1 lists the location of the positive and negative peaks within the 128 position buffer for each of the three phases. The magnitude of the sinewave to be subtracted is determined twice per mains cycle by looking at the positive and negative peaks of the 50Hz current. The three phases are independent and so six magnitude measurements are made during each mains cycle. The sine table is pre-calculated to increase the calculation speed. The sinewave for each phase was scaled appropriately and subtracted from the load current measurement, giving the harmonic current. The compensating current signal is the negative of the harmonic current and is transmitted from the serial port as an integer value.

Table 6.1 Location of the positive and negative peaks of 50Hz currents in the data buffers after low pass filtering.

Phase	Positive Peak	Negative Peak
Red	117	53
Yellow	32	96
Blue	75	11

6.3.3.3 Fast Fourier Transform

Fast Fourier Transform (FFT) based harmonic isolation is very different from the other techniques investigated. The majority of calculations are performed in the main DSP program rather than in an Interrupt Service Routine. The ISRs are still used to acquire a buffer 128 samples long for each of the three phases, and to transmit the compensating current signal.

The forward FFT routine (Tessarolo, 1991), `ffft_rl`, takes 128 sample points in the time domain and converts these to 64 points in the complex frequency domain. The complex frequency values are stored as real and imaginary pairs, giving 128 points in each output vector. The real parts are stored in the first half of the buffer with the DC component being the first term. The imaginary parts are stored in reverse order from the end of the buffer, with the imaginary part of the 50Hz component being at the very end of the output buffer, and this is shown diagrammatically in Figure 6.9.

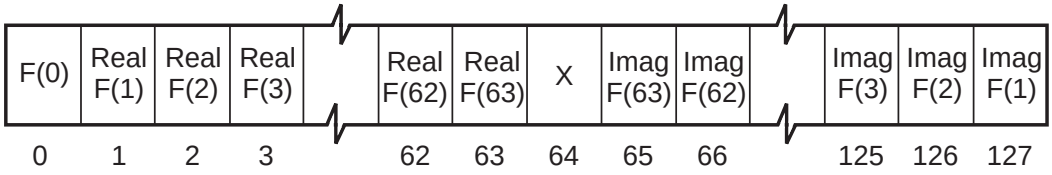

Figure 6.9 Structure of the output buffer from the Fast Fourier Transform routine.

The harmonic current is isolated by removing the DC and fundamental frequency components. This was achieved by setting the contents of locations 0, 1 and 127 in the FFT output buffer to zero. The Inverse Fast Fourier Transform (IFFT) routine (Mazzocco, 1992), `ifft_rl`, takes a complex frequency vector and returns the corresponding real time domain sequence. The input data structure for this routine is exactly the same as the output of the forward FFT shown in Figure 6.9. The IFFT was performed on this altered buffer to create the time domain compensating current signal.

The possibility exists with this technique for more sophisticated compensations such as selective harmonic cancellation to take place. More than one frequency component can be set to zero to prevent the cancellation of these harmonics, or scaling can be applied so there is incomplete harmonic elimination. An investigation into selective harmonic cancellation was performed and the results are presented in §8.6.

6.4 PC Display Software

An IBM compatible computer was used to load software into the Spirit-30 development board's dual port RAM and to display the results of harmonic compensation by the DSP based active filter controller. The software that ran on the PC performed three main tasks:

- Loading of software into the DSP
- Display of the THD and RMS for the load and supply currents
- Display of current waveforms

6.4.1 Control of the DSP

The Spirit-30 development board's ISA interface allowed the PC to control the flow of programs on the DSP. A section of the PC software is shown in Figure 6.10. The first line loads the executable file that contains the DSP program into dual port RAM. Once the DSP program is loaded, the reset line on the DSP was toggled to start the DSP executing the new software. Before a DSP variable in dual port RAM could be read by the PC with the `dsp_up_int_array` command, the memory location of that variable was found with the `get_laddr` command. These two commands, along with `dsp_dl_int_array`, give the PC the ability to read and write any global variable in the DSP's external memory.

```
status = dsp_dl_exec(filename);            // Load the COFF file
dsp_reset();                               // Reset the DSP
          ⋮
dspPC_CMD = get_laddr("_pc_cmd");          // Get address of 'pc_cmd'
          ⋮
dsp_up_int_array(dspPC_CMD, 1, &pc_cmd);   // Get contents of 'pc_cmd'
```

Figure 6.10 Section of PC software responsible for loading the DSP code and reading variables on the DSP.

6.4.2 Active Filter Instrumentation

A special variable, `pc_cmd`, was used for 'handshaking' between the DSP and the PC. By writing certain values from one processor and reading and checking them from the other the two programs could ensure that they were both in a known state. Reading `pc_cmd` by the DSP to check the status of the PC was shown in Figure 6.1 as "Read PC Status". Setting

`pc_cmd` to the value `BUFFER_CONV` by the DSP told the PC that a complete cycle of data was ready for retrieval and display on the PC's monitor.

The binary format of floating point numbers is different between the PC's IEEE format and the 'C30 specific floating point format (Texas Instruments, 1994). Texas Instruments have used the non-standard floating point format to obtain increased efficiency in their hardware implementation of the TMS320C30. The floating point numbers read from the DSP were converted to IEEE format by the PC before they were displayed or used in further calculations. The 32 bit unsigned and signed integer formats of both processors are identical and no conversions were necessary.

The information retrieved from the DSP consisted of the current RMS values, the current THDs and the raw samples of the last mains cycle for each phase. This information was then displayed on the PC screen so comparisons could be made among the different harmonic isolation techniques.

6.4.3 Waveform Plotting

The PC provided excellent graphing facilities through the BGI (Borland Graphics Interface) graphing routines supplied with the Borland Turbo C++ compiler (Borland, 1992). These routines enabled the PC to display four channels of sampled data from the MAX186 ADC in the time domain and the frequency domain. The graphics display showed the version number and name of the DSP software being executed, along with the actual date and time.

Figure 6.11(a) shows the PC software configured as a three phase oscilloscope and spectrum analyser. The traces in the top left graph show the time domain representation of one cycle of load current for all three phases. The frequency spectrum of the three phases was drawn inside a frame. Going clockwise from the top right, the spectra are for Red Phase, Blue Phase and Yellow Phase. The actual supply current was measured with the MAX186 to determine the effect of noise on the quality of compensation. The ideal supply current was calculated mathematically by the DSP from the compensating and load current signal. The actual supply current was the difference between a continuous load current and a sampled compensating current, whereas the ideal supply current was the difference between sampled load and supply currents.

The data shown in the screen capture in Figure 6.11(b) was used for the comparison of harmonic isolation techniques. The traces in the top left box are the load current for the Red Phase and the two supply currents, 'actual' and 'ideal'. The frequency spectrum, RMS current

and THD for the load current, ideal supply current and actual supply current were all shown on the same screen in separate frames. “R THD” and “R RMS” are the THD and RMS current for the Red Phase load current, “S THD” and “S RMS” are for the actual supply current and “C THD” and “C RMS” are the THD and RMS values of the calculated supply current.

The waveforms were downloaded from the DSP in an integer format so no translation of floating point data was necessary. The DSP stored the integer readings from the ADC directly for the PC software and converted the frequency spectrum to an integer format. Scaling is performed by the PC because the screen resolution is less than that of the sampled data. The data was updated by the DSP 50 times per second, but the PC was not able to download and display all of the information in the 20ms provided and simply updated the screen as fast as possible. A screen resolution of 640 pixels wide by 350 pixels high with 16 colours was selected, as this gave enough resolution to show detail in the waveforms without significantly slowing down the computer. The graphics driver had two screen buffers which allowed ‘double buffering’ of the graphics. This is where one screen buffer is displayed on the monitor while the other is redrawn or updated. Each time the screen was updated, the two display buffers were swapped, giving the appearance of an instantaneous update with no screen flicker.

6.4.4 Program Flow

The operation of the PC software has been described in the previous sections. The flow diagram in Figure 6.12 shows the sequence of operation that occurs when the PC software is run and is broken into two sections: the code that is run to initialise the DSP and the PC software, and the code that then runs to provide graphing and instrumentation.

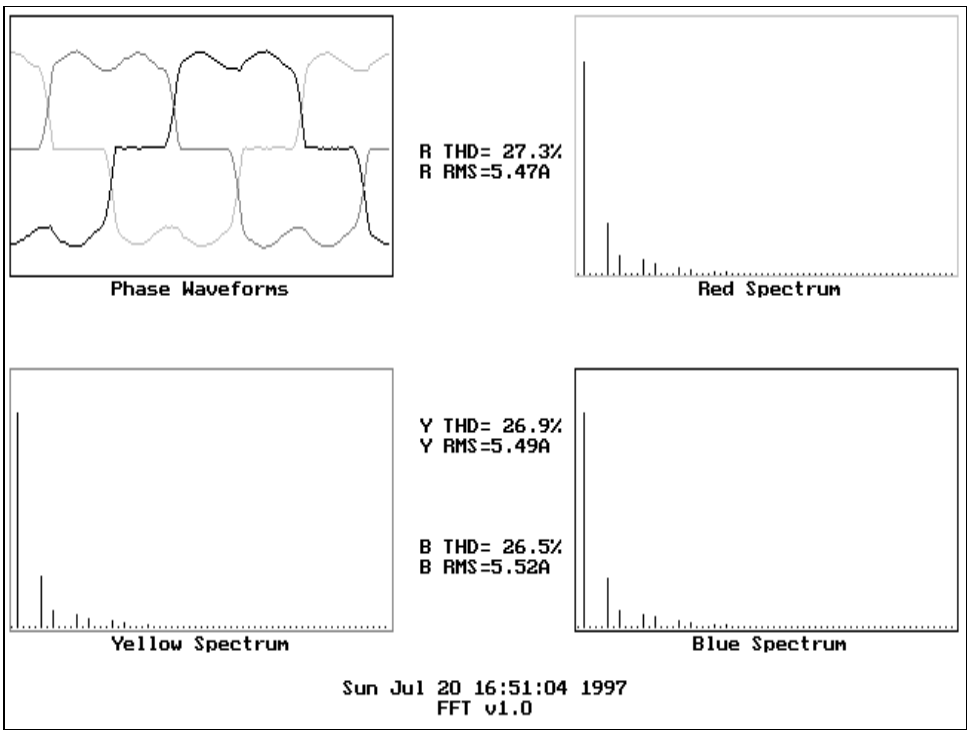

(a)

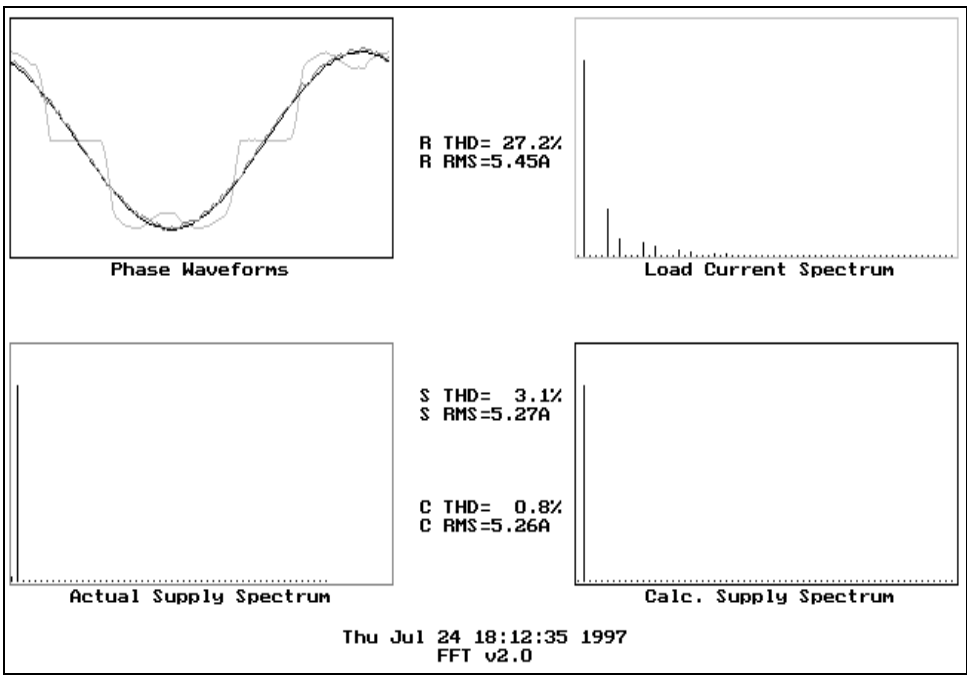

(b)

Figure 6.11 Screen captures from the PC software showing (a) three phase currents in the time and frequency domain and (b) the load and compensated supply currents for the Red Phase.

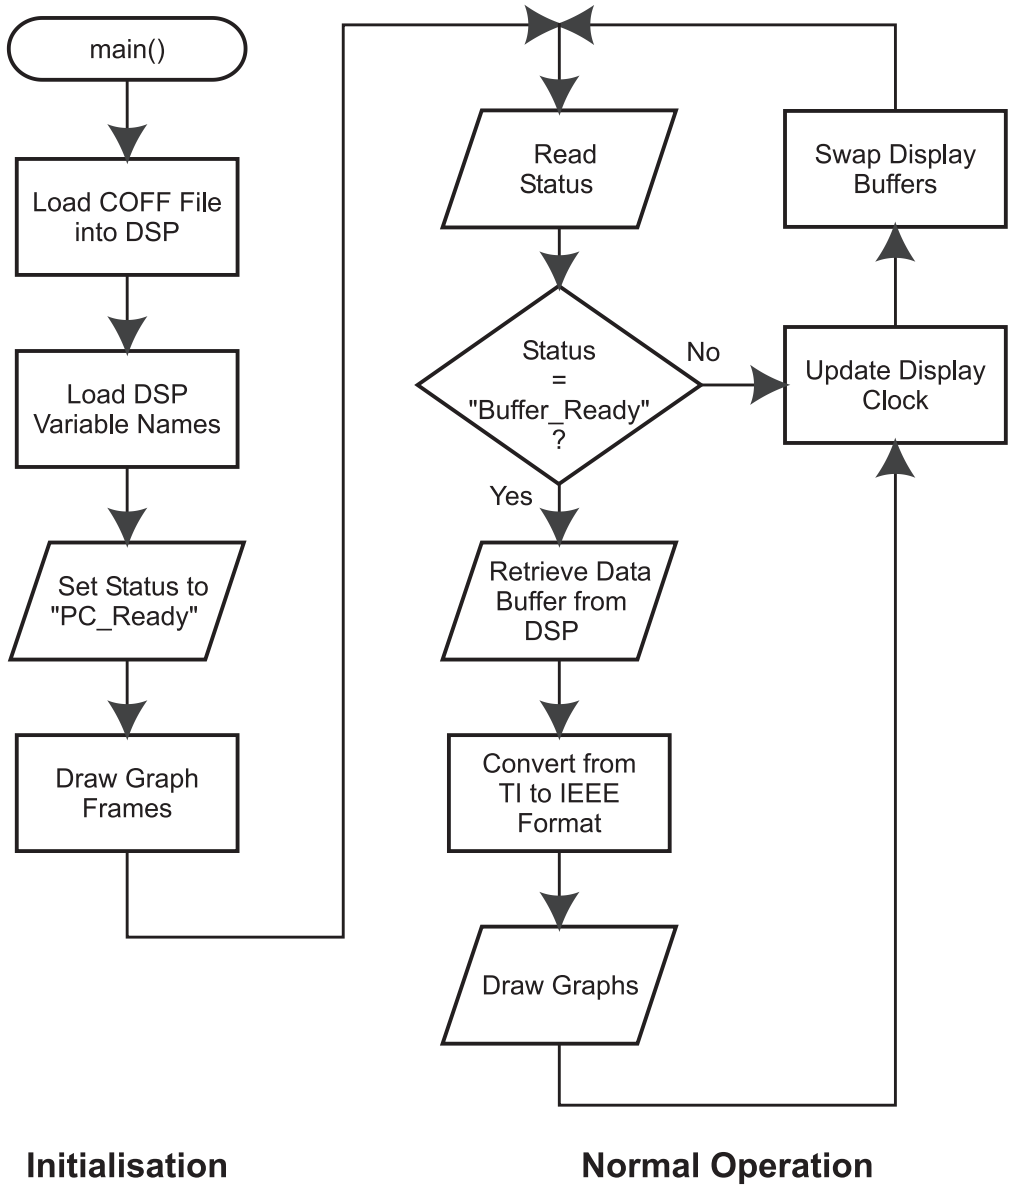

Figure 6.12 Flow diagram of software execution on the PC control computer.

The first task of the PC software was to load the DSP software into the dual port RAM and start the DSP executing that software. Once the software was successfully loaded, the memory locations of the variables to be read and written by the PC were obtained. Handshaking between the DSP and the PC ensured that both pieces of software were operating correctly and at a known ‘holding point’. Frames around the four plots on the screen were drawn into both screen buffers which meant that only the internal area needed to be updated.

Once the initialisation was complete, the PC polled the DSP to check whether a new buffer of data was ready for downloading. If the PC was left waiting for data the clock was

continually updated to show that the PC was operating. This helped distinguish between ‘crashes’ in the PC software and in the DSP software. When new data was ready, the RMS and THD values were converted from Texas Instruments floating point format to the IEEE format and displayed by the PC. The graphs and updated values were drawn onto the ‘invisible’ display buffer along with the new date and time. Once the display update was complete the ‘invisible’ and ‘visible’ display buffers are swapped to show the new data. This process was repeated until the user exited the PC software.

6.5 Summary

The individual elements of the active filter software system have been discussed in this section, with attention to the structure of the software. This design allowed the different harmonic isolation methods to be evaluated experimentally using essentially the same software. One area in which the techniques differed was the digital filtering used to isolate the harmonic components of the load current. The next chapter discusses the theory and implementation of the digital filters used in this work and shows their frequency characteristics.

Chapter 7
Digital Filter Design

Three of the harmonic isolation methods evaluated experimentally used digital filters, both Finite Impulse Response and Infinite Impulse Response. This chapter describes the design of these filters and shows the theoretical responses of each along with an experimental validation. The first section presents a brief look at the theory of digital filters.

7.1 Digital Filter Theory

The mathematics of difference equations allows the frequency characteristics of waveforms to be manipulated in the time domain (Chassaing, 1992). The two classes of digital filters are transversal and recursive. Transversal filters are often called Finite Impulse Response (FIR) filters and recursive filters are called Infinite Impulse Response (IIR) filters. Table 7.1 shows significant differences between FIR and IIR filters and the next two sections will discuss these filters in more detail.

Table 7.1 Summary of differences between FIR and IIR digital filters.

	FIR (Transversal)	**IIR (Recursive)**
Impulse Response	Finite	Infinite
Poles	No finite poles	Finite poles
Taps Required	Several hundred	10 – 20
Stability	Always stable	May be unstable
Phase	Can be exactly linear	Cannot be linear
All-Pass Filters	Cannot realise	Can realise
Limit Cycle Oscillations	Free from oscillations	Can oscillate due to round off and quantisation noise.

7.1.1 Finite Impulse Response Filters

A Finite Impulse Response (FIR) filter is designed so that $N+1$ coefficients, $h(0), h(1), \ldots, h(N)$, give a transfer function $H\left(e^{j\omega}\right)$ that approximates a desired frequency characteristic $H_d\left(e^{j\omega}\right)$. All other impulse response coefficients, which are also called "taps", will be zero. An important property of FIR filters is that they can be designed to have a linear phase characteristic (Poularikas *et al.*, 1993).

Two methods exist for implementing FIR filters: direct evaluation of the convolution sum and the use of the overlap-save Discrete Fourier Transform (Poularikas *et al.*, 1993). The direct form implementation of the FIR filter will be discussed in this thesis. This method performs a convolution between the input vector (the present and delayed values of the input) and the weighting vector of $N+1$ coefficients. The 'convolution sum' is expressed in Eqn. (7.1) with $h[0] \ldots h[N]$ representing the impulse response and $x[0] \ldots x[N]$ the input vector.

$$y[n] = h[0]\,x[n] + h[1]x[n-1] + \ldots + h[N]\,x[n-N] \tag{7.1}$$

This expression is shown diagrammatically in Figure 7.1 as a series of delay blocks (indicated by "z^{-1}"), scaling and adders (+). When a new input value is entered into the filter it is placed at the beginning of the input sequence and all existing values move back one time position. The input samples pushed past the end of the vector are discarded since the weighting vector is assumed to be zero at this point. For example, when a new $x[n]$ is formed, the previous value of $x[n]$ becomes $x[n-1]$, the previous value of $x[n-1]$ becomes $x[n-2]$ and so forth. This ripples through until there are no more delay blocks and the old value of $x[n-N]$ is discarded. The filter coefficients remain constant for all filter calculations.

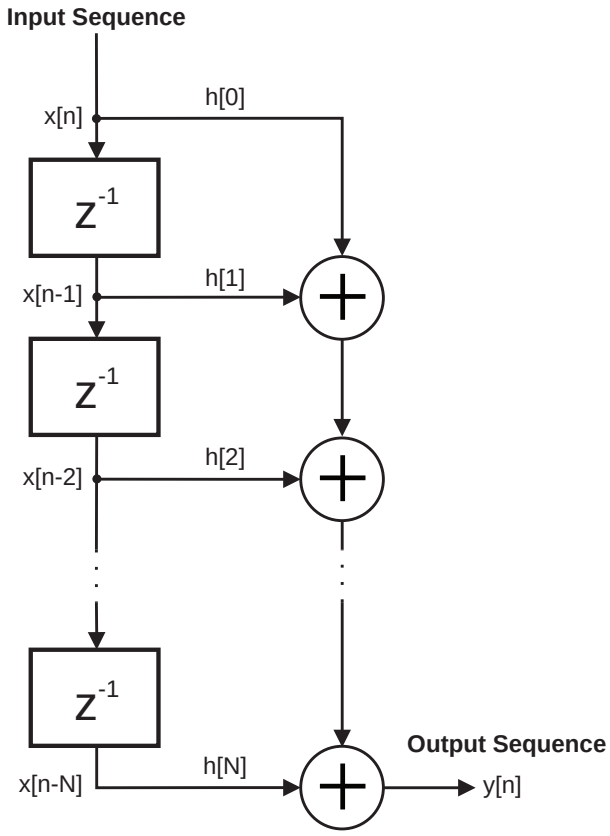

Figure 7.1 Block diagram of an FIR filter.

7.1.2 Infinite Impulse Response Filters

The transfer function of an IIR filter has both poles and zeros, with the output depending upon previous inputs and outputs. These filters generally take less computation than an FIR filter with a similar frequency response. IIR filters do have the disadvantage of being sensitive to coefficient quantisation (Chambers, Tantaratana and Bomar, 1996).

IIR filters are commonly designed by taking an analogue filter with the desired frequency response and then forcing the digital filter to have the same response. Two commonly used methods are:

- **Impulse Invariant**

  The impulse response of the digital filter is the sampled impulse response of the analogue filter. This sampling is described in Eqn. (7.2), where F_S is the sampling frequency of the digital filter.

$$h(n) = h_a(nT),\ n = \frac{1}{F_S} \tag{7.2}$$

- **Bilinear Transformation**

  The analogue transfer function, $H_a(s)$, is transformed into the digital transfer function, $H(z)$, with the expression given in Eqn. (7.3). "Prewarping" of the analogue frequency characteristic is required because the bilinear transformation is non-linear (Poularikas *et al.*, 1993).

$$H(z) = H_a(s)\Big|_{s=\frac{2}{T}\frac{1-z^{-1}}{1+z^{-1}}} \tag{7.3}$$

The equation for a "Direct Form I" IIR filter is given by Eqn. (7.4) and shown diagrammatically in Figure 7.2. This IIR filter has two delay chains, as opposed to a single chain in the FIR filter of Figure 7.1. The input values are delayed by the chain on the left-hand side of the figure and the output values are delayed by the right-hand chain.

$$\begin{aligned} y[n] = a_1\, y[n-1] + \ldots + a_N\, y[n-N] + \\ b_0\, x[n] + b_1 x[n-1] + \ldots + b_M\, x[n-M] \end{aligned} \tag{7.4}$$

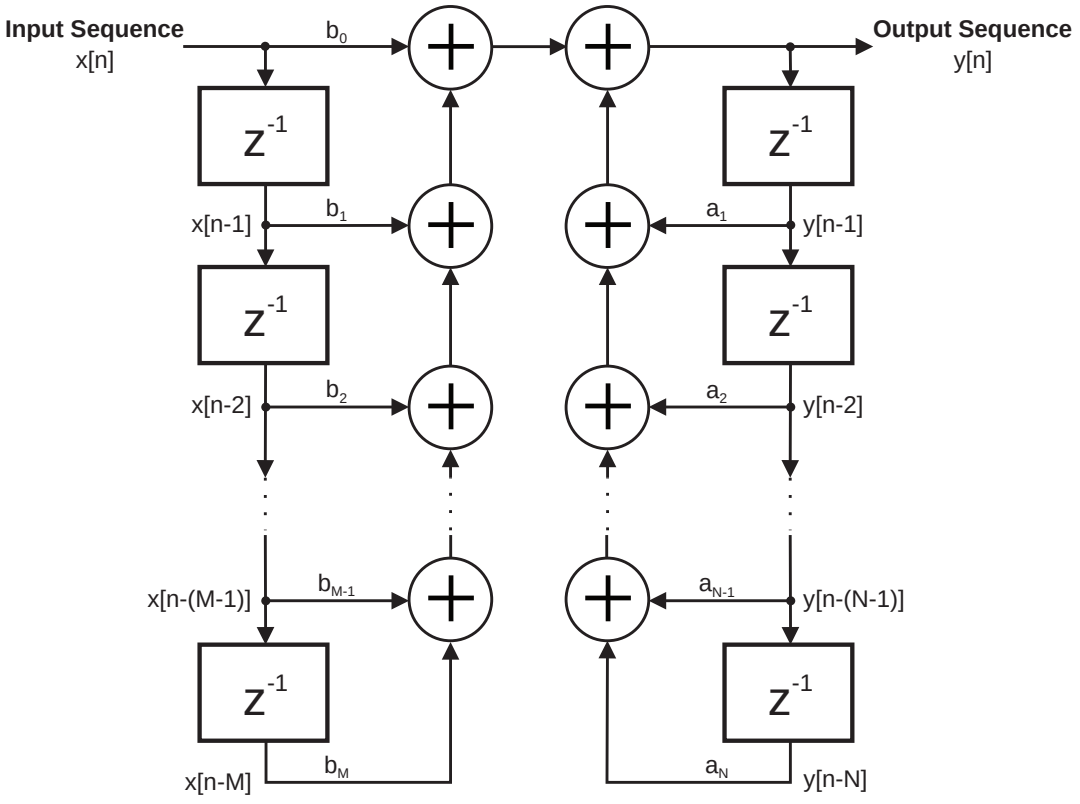

Figure 7.2 Direct Form I realisation of a recursive filter.

The IIR filter can be rewritten using the two equations given by Eqn. (7.5) as long as there are sufficient taps to accommodate *M* and *N*. Any unused taps have their value set to zero. This form is more convenient, has reduced storage needs and is called the "Direct Form II" realisation. A single chain of delay elements is used, as shown in Figure 7.3. The delay chain is in the centre of the figure with the multiplicative taps branching off from each side.

$$\begin{aligned} d[n] &= x[n] + a_1\, d[n-1] + \ldots + a_N d[n-N] \\ y[n] &= b_0\, d[n] + b_1\, d[n-1] + \ldots b_N\, d[n-N] \end{aligned} \tag{7.5}$$

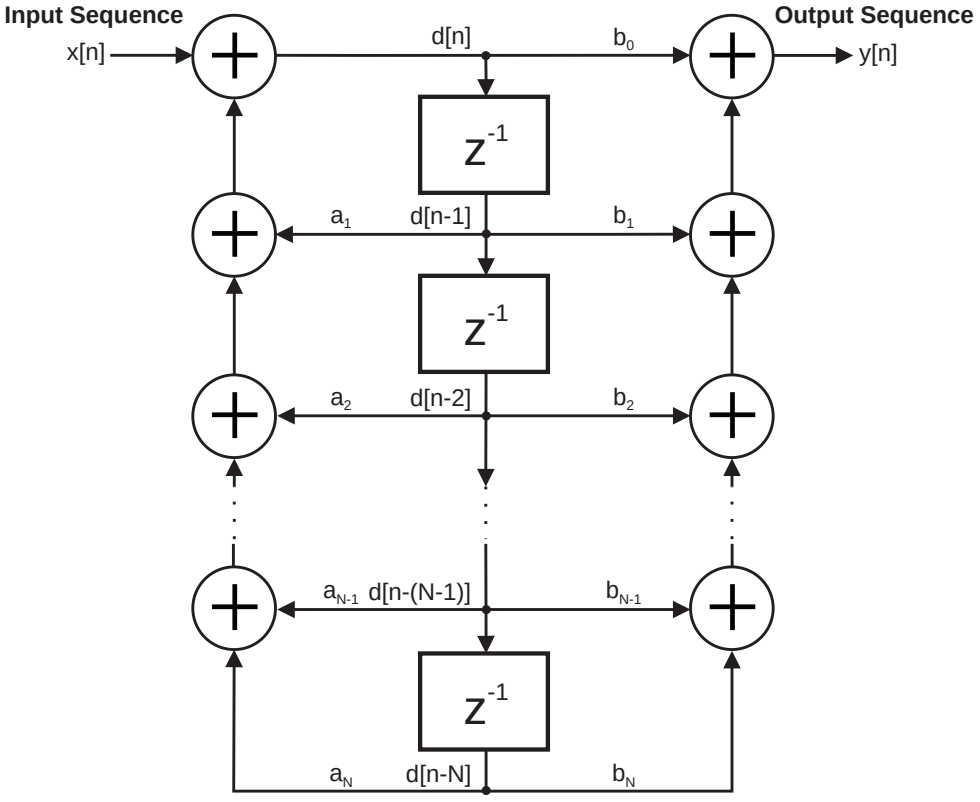

Figure 7.3 Direct Form II realisation of an IIR filter.

IIR filters cannot achieve the linear phase that an FIR filter can and suffer from "finite word length effects". There are four types of finite word length effects that affect IIR filters: round-off noise, coefficient quantisation error, adder overflow limit cycles and multiplier round-off limit cycles (Poularikas *et al.*, 1993).

The two types of limit cycle require recursion to occur and therefore are not present in FIR filters. Direct Form I and Direct Form II IIR filters generally have poor performance due

to finite word length effects (Chambers *et al.*, 1996). These effects are reduced by breaking the filter down into a cascade of second order IIR filters called ‘biquads’. Factoring or expanding the original transfer function significantly reduces quantisation error and round-off noise (Poularikas *et al.*, 1993). Figure 7.4(a) shows the structure for *r* cascaded blocks. The structure of a fourth order IIR filter made with two Direct Form II sections is shown in Figure 7.4(b).

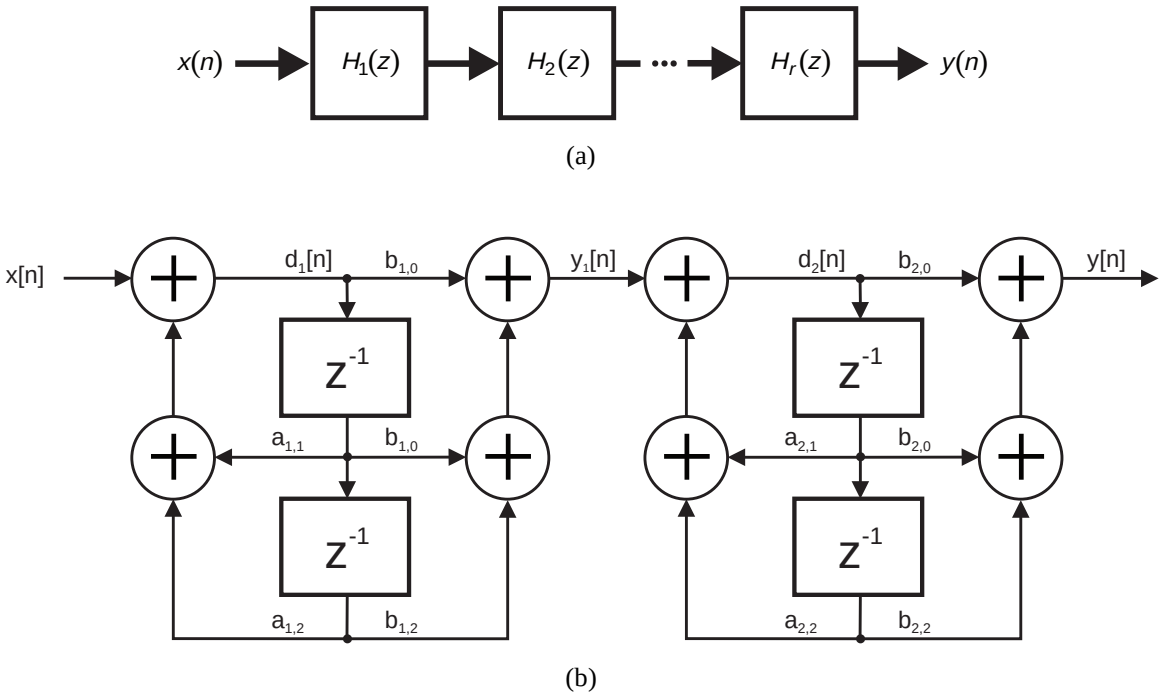

Figure 7.4 Cascade IIR structure (a) and a fourth-order IIR filter with two Direct Form II sections (b).

7.2 Calculation of Filter Coefficients

7.2.1 Filter Requirements

The high pass filter had to remove the 50Hz fundamental current component without affecting the harmonic currents, and a filter with a 3dB point at 100Hz satisfied these requirements. The low pass filter had to pass the 50Hz fundamental current while removing all harmonic values from the third onwards. The 3dB attenuation frequency was set at 90Hz, which was a trade off between no attenuation at 50Hz and complete rejection at 150Hz. The IIR notch filter has a non linear phase response, with the phase following that of an analogue

notch filter. A first order filter was selected to reduce the phase error, as this would have lead to incorrect harmonic elimination. The 3dB bandwidth affects the depth of the notch, but by using a Sample Rate Multiplier to alter the sampling frequency, the centre of the notch tracks the mains frequency. This allows the use of a narrow bandwidth filter to give increased rejection of the 50Hz current signal.

7.2.2 Filter Calculation

The “Signal Processing Toolbox”, an add on for MATLAB (Mathworks, 1996a), was used to calculate the coefficients of each digital filter used in the active filter controller. Four filters were created:

- 128 Tap FIR High Pass Filter with -3dB at 100Hz
- 256 Tap FIR High Pass Filter with -3dB at 100Hz
- 256 Tap FIR Low Pass Filter with -3db at 90Hz
- 3 Tap IIR Notch Filter with -3dB at 48Hz and 52Hz

Hamming windows were used on each of the FIR filters to reduce the amplitude of the sidelobes, as this has previously been shown to give good results (Chassaing, 1992). The frequency characteristics of these filters are shown in Figure 7.5. The two high pass filters, Figure 7.5(a) and Figure 7.5(b), show how doubling the number of taps increases the steepness of the transition from stopband to passband.

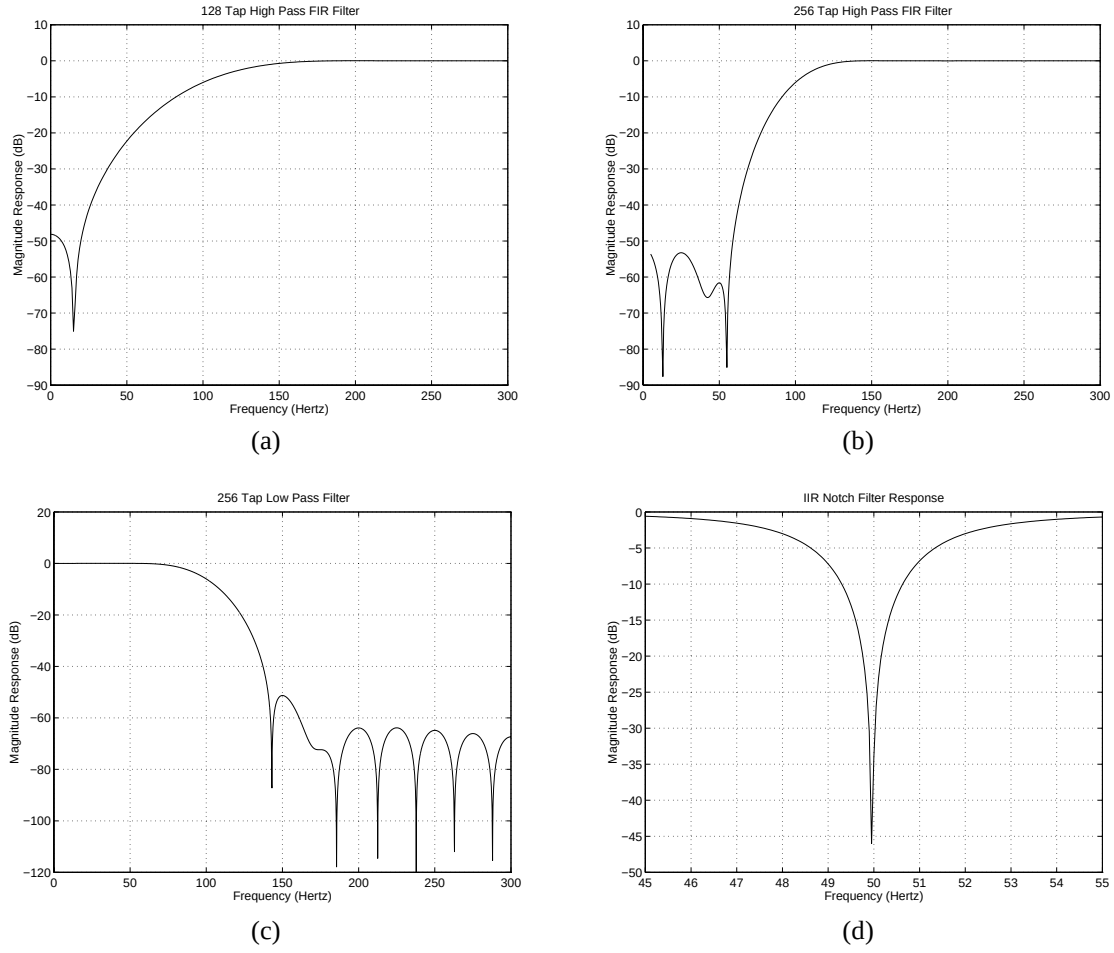

Figure 7.5 Frequency domain magnitude characteristics of the four digital filters used in the active filter controller.

7.3 DSP Implementation of the Filters

The *C3x User's Guide* (Texas Instruments, 1994) contains assembly language routines for FIR and IIR filtering that are extremely efficient. The next two sections describe the operation of these routines.

7.3.1 Finite Impulse Response

The TMS320C30 has two features that assist the development of an FIR filter. These are parallel multiply-add operations and circular addressing. The parallel multiply-add performs the 'Multiply and Accumulate' (MAC) function of other DSPs but is more flexible in its operation (Texas Instruments, 1994).

Many algorithms, such as convolution and correlation, require the use of a circular buffer held in memory. The circular buffer implements a sliding window that contains the most recent samples, with the new information overwriting the oldest. Circular addressing 'loops' an auxiliary register around a fixed piece of memory when the register is incremented, rather than increasing linearly. The algorithm used for circular addressing is given in Figure 7.6 with the condition that the step size (`step`) must be less than or equal to the block size (`BK`), and is treated as an unsigned integer. The current location in the buffer is given by '`index`'.

```
If 0 ≤ (index + step) < BK Then
      index = index + step.
Else If (index + step) ≥ BK Then
      index = index + step – BK.
Else If (index + step) < 0 Then
      index = index + step + BK.
```

Figure 7.6 Circular addressing algorithm, modified from the *C3x User's Guide* (Texas Instruments, 1994).

If the proposed destination is within the buffer, the step is applied to the current index. If the new position would exceed the limits of the buffer, the step is added and the buffer length subtracted. Conversely, if a step would take the index back past the first point in the buffer, the buffer length is added to the result.

Circular addressing is set up by initialising the block size register, `BK`, to the length of the buffer, *N*. The start of the buffer that holds the input signal *x* must start from a memory location whose address is a multiple of the smallest power of two that is greater than *N*, the length of the filter. For example, if $N = 24$, the first address for *x* should be a multiple of 32 (the lowest five bits of the beginning address should be 0).

The assembly code shown in Figure 7.7 implements an FIR filter on the 'C30 DSP. Five registers are used in this routine:

- R0 Result of multiplication of a delayed input value and the corresponding tap
- R2 Sum of all tap and weight multiplications
- AR1 Pointer to the weight sequence
- AR2 Pointer to the input sequence
- RC Repeat Counter that gives the length of the FIR filter

AR1 and AR2 are post-incremented by the '++(1)' suffix and this steps the DSP through the taps without any processor overhead. '*AR1++(1)%' increments the AR1 auxiliary register with circular addressing because of the trailing '%' character. The 'block size' of the circular buffer is held in the 'BK' register and when the register is incremented past this point it is set back to the start of the circular buffer. 'RPTS RC' repeats the instruction on the line following RC times without any overhead for looping. In this example, the next instruction takes two lines to write because it is a parallel instruction, indicated by the '||' at the start of the second line. The DSP performs the multiplication and the addition simultaneously. The addition uses the value of R0 at the start of the instruction while the multiplication is generating a new value of R0. This requires a single multiplication to be performed ahead of the repeated parallel instruction and a single addition to be performed afterwards to take into account this delay.

```
   LDF    0.0, R2                     ; Initialise R2
   MPYF3  *AR2++(1), *AR1++(1)%, R0   ; Initial Multiply

   RPTS   RC                          ; Loop N-2 times
   MPYF3  *AR2++(1), *AR1++(1)%, R0   ; q[n] = x[n] * h[n]
|| ADDF3   R0, R2, R2                 ; y[n] = q[0] + ... + q[n]

   ADDF3  R0, R2, R0                  ; Add the last product
```

Figure 7.7 Assembly language listing for a FIR filter implemented on a TMS320C30 DSP.

7.3.2 Infinite Impulse Response

Figure 7.8 shows the assembly code implementation of a second order IIR filter using a single biquad. The register usage is:

- R2 Input Sample
- AR0 Address of filter coefficients
- AR1 Address of delay mode values.
- AR2 Address where filter result is to be stored
- BK Block size of three
- R0, R1 Temporary floating point registers

This filter implementation uses the Direct Form II structure shown in Figure 7.3. The first four sections of code in Figure 7.8 compute the individual terms required by the filter.

The fifth, and final section, adds each of the scaled delay values together to give the output of the filter.

```
   MPYF3  *AR0, *AR1, R0        ; a2*d[n-2] ->R0
   MPYF3  *++AR0, *AR1--%, R1   ; b2*d[n-2] ->R1

   MPYF3  *++AR0, *AR1, R0      ; a1*d[n-1] ->R0
|| ADDF3  R0, R2, R2            ; a2*d[n-2] + x(n) ->R2

   MPYF3  *++AR0, *AR1--%, R0   ; b1*d[n-1] ->R0
|| ADDF3  R0, R2, R2            ; a1*d[n-1] + a2*d[n-2] + x[n] ->R2

   MPYF3  *++AR0, R2, R2        ; b0*d[n] ->R2
|| STF    R2, *AR1++%           ; Store d[n] and point to d[n-1]
   STI    AR1, @DEL POS         ; Update the DELAY position

   ADDF   R0, R2                ; b1*d[n-1] + b0*d[n] -> R2
   ADDF3  R1, R2, R0            ; b2*d[n-2] + b1*d[n-1] + b0*d[n] ->R0
   STF    R0, *AR2              ; Return result
```

Figure 7.8 IIR Filter routine that calculates one biquad.

The IIR coefficients are stored in a single piece of memory in a specific order. This allows the pre-increment and post-increment operators (eg. *++AR0 and *AR1++%) to access the coefficients in the most efficient manner. Table 7.2 lists the coefficients in the order in which they are stored in memory.

Table 7.2 Memory Organisation of the filter coefficients.

Position	Weight
1	a2
2	b2
3	a1
4	b1
5	b0

7.4 Swept Sine Filter Testing

A Hewlett-Packard HP35665A Dynamic Signal Analyser (DSA) generated a swept sine signal that ranged from 5Hz to 1kHz. The active filter controller was programmed so the output of the digital filter was turned into an analogue voltage by an external digital to

analogue converter. This allowed the DSA to plot the magnitude characteristic of each digital filter.

The magnitude plots generated by the DSA were compared with the magnitude characteristics generated by MATLAB to verify that the implementation of each digital filter was correct. To demonstrate this comparison, the MATLAB and DSA plots of the 256 tap FIR high pass filter and three tap IIR notch filter responses are shown in Figure 7.9. It can be seen that the magnitude characteristics of the simulation and the DSP digital filter are almost identical. This validates the DSP implementation of the digital filters and removes the need to test every digital filter before use.

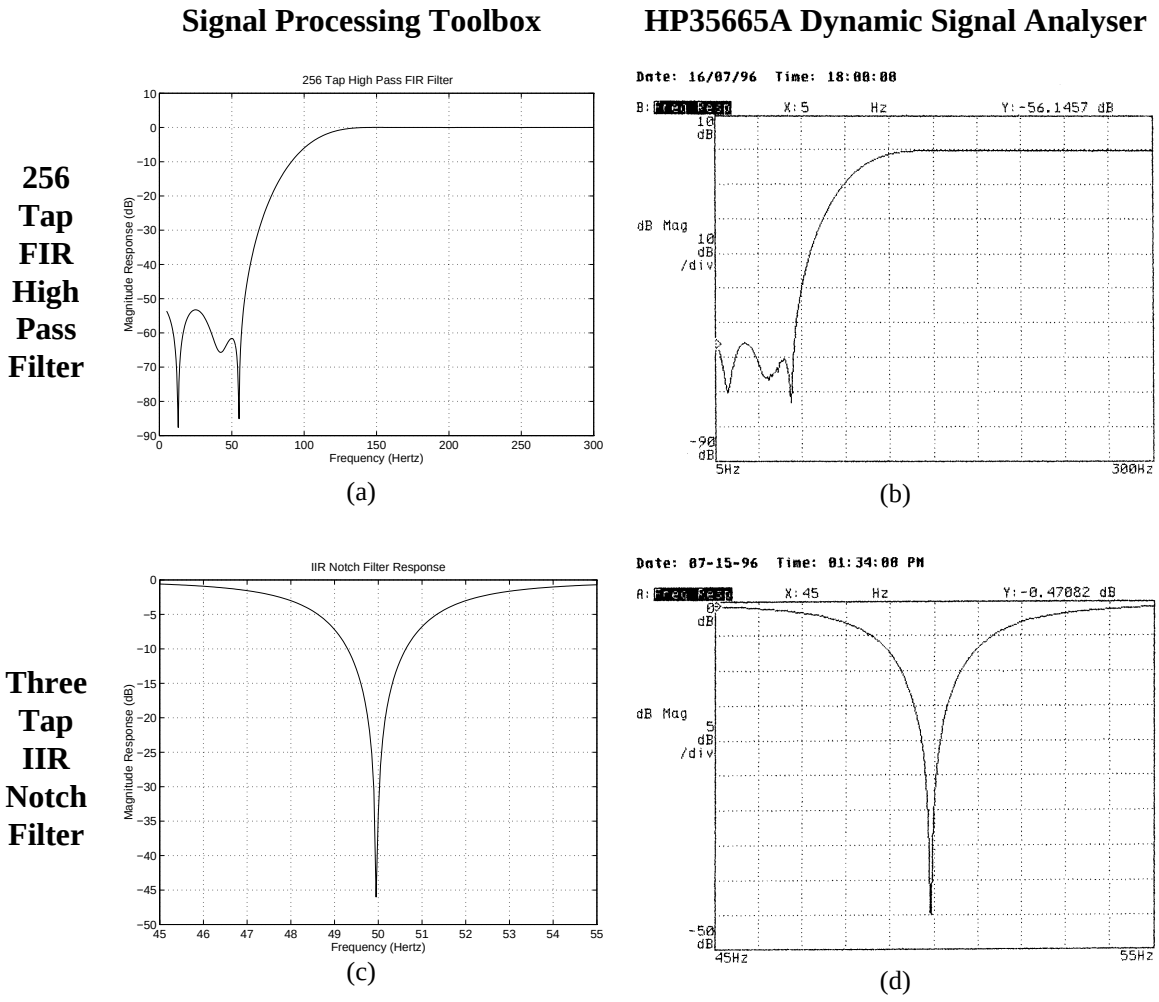

Figure 7.9 Frequency domain plots of the FIR filter characteristic (a) and (b), and the IIR filter characteristic (c) and (d).

7.5 Summary

FIR and IIR digital filters provide an efficient means of filtering signals. Each has distinct characteristics that affect its suitability for a particular application. FIR filters have an ideal linear phase response, but require a large number of coefficients, whereas IIR filters are computationally efficient but have less than ideal transient and phase responses.

Filter designs from MATLAB's Signal Processing Toolbox have been tested by swept sine analysis of the digital filters implemented in the TMS320C30 DSP. The results shown in this chapter, particularly those of Figure 7.9, verify that the digital filters have been correctly implemented in the DSP software. The digital filters presented in Figure 7.5 were used in the experimental system, with the result of these tests presented in the next chapter.

Chapter 8
Experimental Evaluation of Harmonic Isolation Techniques

This chapter presents the results of experimental implementations of four harmonic isolation methods: Notch Filtering, High Pass Filtering, Sinusoidal Subtraction and the Fast Fourier Transform. The relative merit of each method was evaluated by looking at the transient performance and steady-state quality of compensation. The computational complexity of each method was measured as this determines the processing power required by the DSP to implement a full-scale active filter.

8.1 Apparatus Used for Testing

A three phase bridge rectifier generated a load current with high harmonic content. Using an inverter to inject current into the mains may have masked any errors introduced by the harmonic isolation method. To prevent this, the active filter controller produced an analogue compensating current signal that was added with an op-amp summer circuit to the load current signal. Figure 8.1 shows this equipment in schematic form.

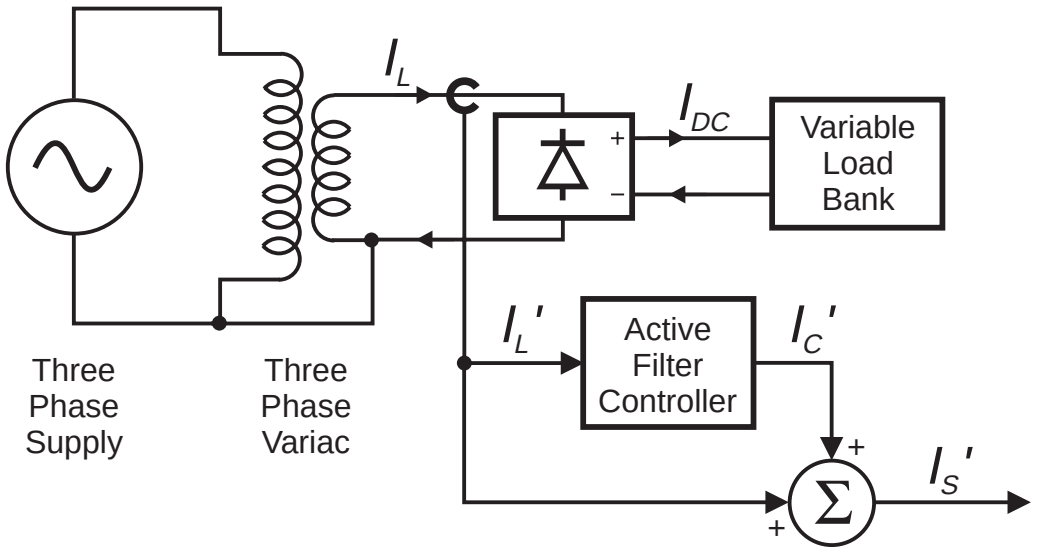

Figure 8.1 Test apparatus for evaluating the harmonic isolation methods.

8.2 Computational Complexity

The computational complexity of the algorithms must be considered when selecting harmonic isolation methods. A complete active filter controller has many tasks to perform, in addition to isolating the harmonic content of the load current. Supervisory calculations must be performed to ensure that the active filter is operating correctly and safely. The DC bus voltage on the capacitors used to supply the voltage source inverter must maintained by regulating the flow of real power into, and out of, the active filter. These calculations, and many others, all place a processing load on the Digital Signal Processor (DSP). The controller is a real-time system and has a fixed amount of time to perform calculations. In a 50Hz power system all cyclic calculations must be completed within the mains cycle period of 20ms.

A Texas Instruments XDS-510 In Circuit Debugger was used to count the clock cycles taken by the DSP to execute each harmonic isolation algorithm. Each clock cycle took 60ns in the 33MHz TMS320C30 DSP used as the controller, and therefore the time required for each algorithm to execute is known. Two measurements were made for each technique, the number of clock cycles required to perform the filtering, and the clock cycles required to call and execute the assembly language filter function from the main C program. This gave an indication of the program overhead resulting from passing the input parameters to the filter routine and preserving the registers used by the filter routines.

The filters were executed from both the high speed internal RAM and the external dual port RAM on the Spirit-30 board to compare the effect of different memory access speeds. The Notch Filtering harmonic isolation method used an IIR filter, while the High Pass Filtering and Sinusoidal Subtraction methods used FIR filters. The total number of instruction cycles per sample can be calculated from the sampling period and the instruction cycle length. For a nominal 50Hz fundamental the sampling period is 156.25µs, which at 60ns per instruction cycle gives a total of 2604 instruction cycles available per sampling interval for all three phases. From this, the number of instruction cycles per phase, including any background processing overhead, is 868.

8.2.1 IIR Notch Filter

Nine clock cycles were required to execute the IIR filter code from internal RAM. Executing from external RAM slowed the filter down remarkably, with the IIR filter taking 23 clock cycles to execute. When the overhead generated by calling an external assembly routine was considered, the total time was 69 clock cycles for the filter running in internal

RAM and 80 clock cycles when run in external RAM. The assembly language calling overhead for internal and external RAM was approximately 60 clock cycles. The number of clock cycles required to perform the filtering is almost insignificant when compared to the overhead, and therefore little advantage is obtained by using the high speed internal RAM. The IIR filter does not load the processor heavily; 7.9% and 9.2% loading when executing from internal and external RAM respectively.

8.2.2 FIR High Pass Filters

The FIR filters have a greater number of taps than the IIR filter and therefore take longer to execute. The assembly code implementation of a FIR filter is extremely efficient, with the number of clock cycles required to perform the filter very close to the number of taps in the filter when memory accesses take one instruction cycle.

When executing from internal RAM, the 128 tap FIR filter took 135 clock cycles and the 256 tap FIR filter took 265 clock cycles to complete. The external RAM once again slowed down the filter calculations, with the 128 tap FIR filter taking 265 clock cycles and the 256 tap FIR filter requiring 521 clock cycles. The memory timing given in the *C3x User's Guide* (Texas Instruments, 1994) is one clock cycle for an internal RAM access and two clock cycles for zero waitstate external memory. The results confirm that external RAM access requires twice as many clock cycles as an internal RAM access.

The C calling overhead for the two FIR filters, in internal and external RAM, was very close to 60 clock cycles. This is similar to the overhead experienced with the IIR filter in §8.2.1. The increased performance of the two FIR filters obtained by using internal RAM was significant, as the number of clock cycles required for performing the filtering is much greater than the calling overhead. In this situation it is prudent to use internal RAM for the FIR filters. The processor load for these routines is much higher, 15.5% for the 128 tap filter and 30.5% for the 256 tap filter when executing from internal RAM.

8.2.3 Sinusoidal Subtraction

The sinusoidal subtraction technique used a 256 tap FIR low pass filter to obtain the magnitude of the fundamental of the load current. No band-pass filter was used for synchronising since the Sample Rate Multiplier (discussed in §4.3) provided this. An advantage of an FIR low pass filter with a tap length equal to double the number of samples per cycle is that the delay through the filter is exactly one cycle for all frequencies. This makes determining the peak of the filtered current waveform much easier as the peak will

always be at a fixed point in the input buffer. The calculations required to scale a sinusoid are negligible in comparison to those required for the low pass filtering. For this reason, the computational complexity is taken to be the same as the 256 tap FIR high pass filter, 30.5%.

8.2.4 Fast Fourier Transform

Harmonic isolation using the Fast Fourier Transform is quite different to the other techniques. Rather than operating on a point by point basis, the FFT method considers the load current cycle by cycle. The ADC digitised the load current one sample at a time and sent the conversion result to the DSP. Once a complete cycle of data was obtained, an FFT was performed on this data. The real and imaginary fundamental components are set to zero and the Inverse Fast Fourier Transform (IFFT) performed on the altered frequency information. The resulting compensating current is then sent to the power inverter one sample at a time, starting at the beginning of the next cycle.

Figure 8.2 shows the stages in the calculation of harmonic compensating currents. Each of the three blocks under a given cycle of current is performed concurrently. The receiving of current samples and the transmission of compensating currents interrupts the compensating current calculations. It can be seen that during any cycle one block (A, B, or C) will be acquiring the load current, another will be calculating the compensating current and another will be injecting the pre-calculated compensating current. The figure is indicative of only one phase and so a three phase active filter controller will perform the calculations three times.

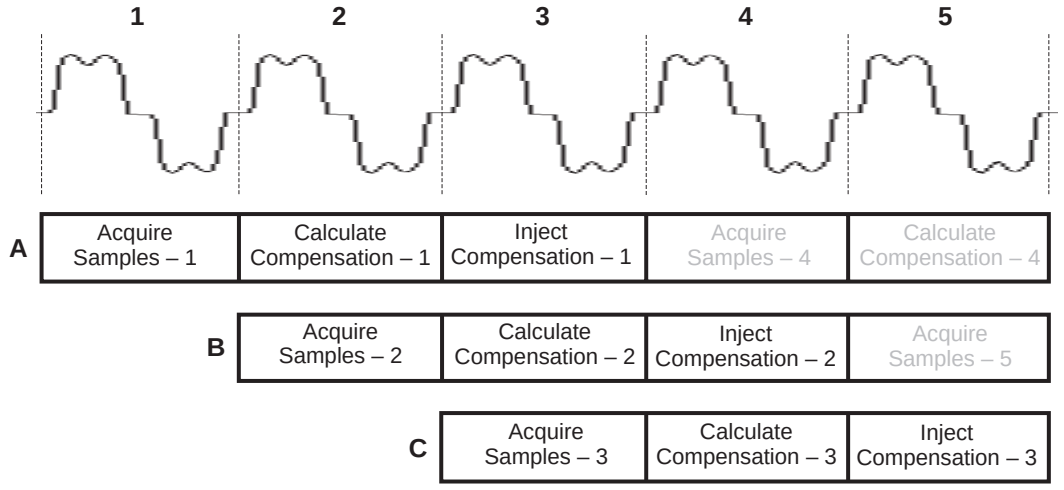

Figure 8.2 The compensation current calculation sequence for the FFT method.

For a single cycle of load current with 128 points, the FFT takes 4956 clock cycles to execute. Setting the fundamental component to zero takes eighteen clock cycles, while the IFFT to generate the compensating current in the time domain takes 6396 clock cycles. This gives a total of 11370 clock cycles per mains cycle, which is effectively 88 clock cycles per sample, giving a processor load of 10.1%. If selective harmonic cancellation is used (as is discussed later in §8.6) the time taken to zero other frequencies will increase the overall number of clock cycles required.

8.2.5 Summary of Computational Complexity

The complexity of each of the routines discussed above can be summarised as the number of clock cycles required per sample to perform the filtering. Table 8.1 shows this summary for each of the methods implemented by the DSP.

Table 8.1 Summary of computational complexity for each harmonic isolation method. The clock cycle count is the lowest possible for each algorithm.

	Clock Cycles per Sample	Clock Cycles per Mains Cycle	Effective Clock Cycles per Sample	Real-Time Processor Loading
128 Tap FIR High Pass	135	—	135	15.5%
256 Tap FIR High Pass	265	—	265	30.5%
IIR Notch	69	—	69	7.9%
Sinusoidal Subtraction	265	—	265	30.5%
Fast Fourier Transform	—	11370	88	10.1%

8.3 Digital Delay Artefacts

It became apparent that the supply current that resulting from all of the harmonic isolation methods examined in this thesis was not perfectly sinusoidal, with small ‘spikes’ appearing in the supply currents. This occurred when the supply current was created with the channel addition function of an oscilloscope and when the addition was performed by the op-amp summer. An example waveform is shown in Figure 8.3, where Channel 1 is the load

current, Channel 2 is the compensating current and Channel M1 is the addition of the two by the oscilloscope.

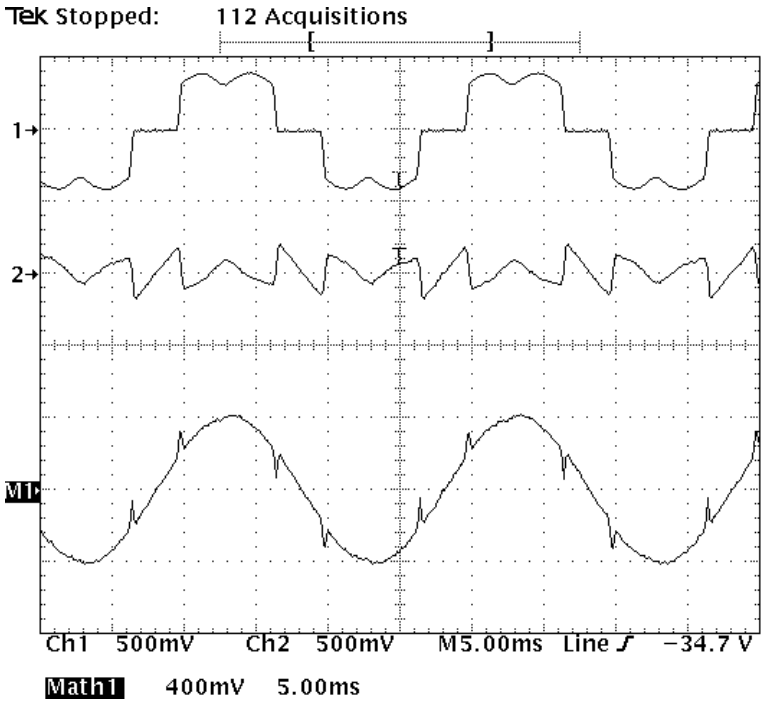

Figure 8.3 Notches that appear in the calculated supply current for an IIR notch filter based active filter.

Further investigation has shown that this ‘notching’ is due to time delays caused by the sampling and processing of the load current signal. When sharp transitions occur, the time difference is enough for errors to form. The error was more pronounced with load current waveforms that were ‘double-humped’ and those with high rates of change of current. The notches were shown to reduce when the sampling frequency of the ADC used to sample the load current increased. Figure 8.4 shows an analogue voltage (Channel 1) and the same voltage when sampled at 6.4kHz (Channel 2). The difference between these two currents is calculated by the oscilloscope (Channel M2). Even though the slope is small, the instantaneous difference between the two currents is quite large.

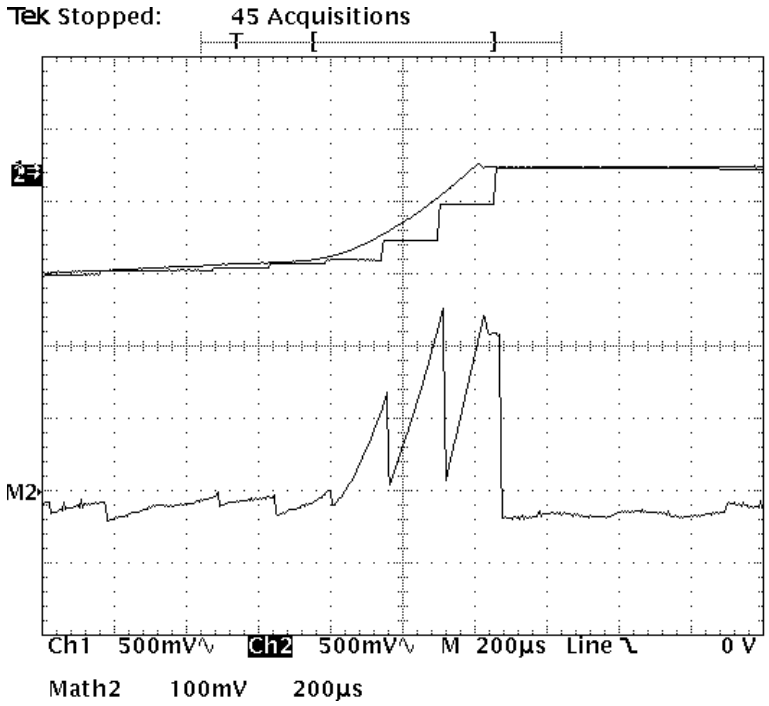

Figure 8.4 The effect of sampling and delaying the load current signal.

8.3.1 Error in the Simulation Model

The simulations presented in Chapter 3 gave extremely good results with no notching of the resulting supply current. This is contrary to the results shown in Figure 8.3 and further examination of the simulation model showed a design flaw. The original model, illustrated in Figure 8.5(a), quantised and sampled the load current for all further calculations, but in reality the load current remains continuous and a sampled and quantised representation of the load current signal is used to determine the compensating current. This sampled and quantised compensating current is added to the continuous load current to give the supply current. The model was changed to reflect this and is shown in Figure 8.5(b).

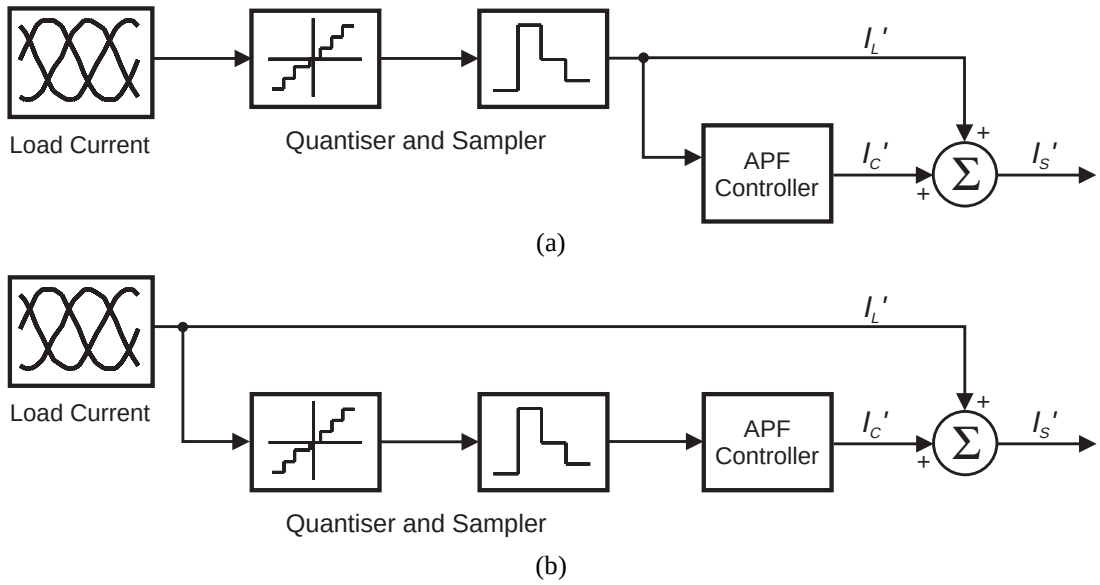

Figure 8.5 The original model used for simulation (a) and the corrected model (b).

8.3.2 Modelling of Effects

The two processes that affected the quality of the compensating current were quantisation and sampling. Twelve bit analogue to digital converters were used, giving 4096 possible current levels. This resolution is adequate since the noise level in many parts of the circuit was greater than the voltage or current equivalent to one bit. Sampling of the load current waveform changed the output far more drastically. With slow sampling, the analogue quantity is 'held' much longer and this gives larger errors. The solution is to sample the load current more often to reduce the hold time.

The quantisation and sampling rate used by the DSP were used to verify the new model. The load current, compensating current and resulting supply current were plotted from the simulation and experimental results. 'Notches' form in the simulated result in Figure 8.6(a) when the new model is used. These occur at the same place as the 'notches' in the oscilloscope capture of Figure 8.6(b).

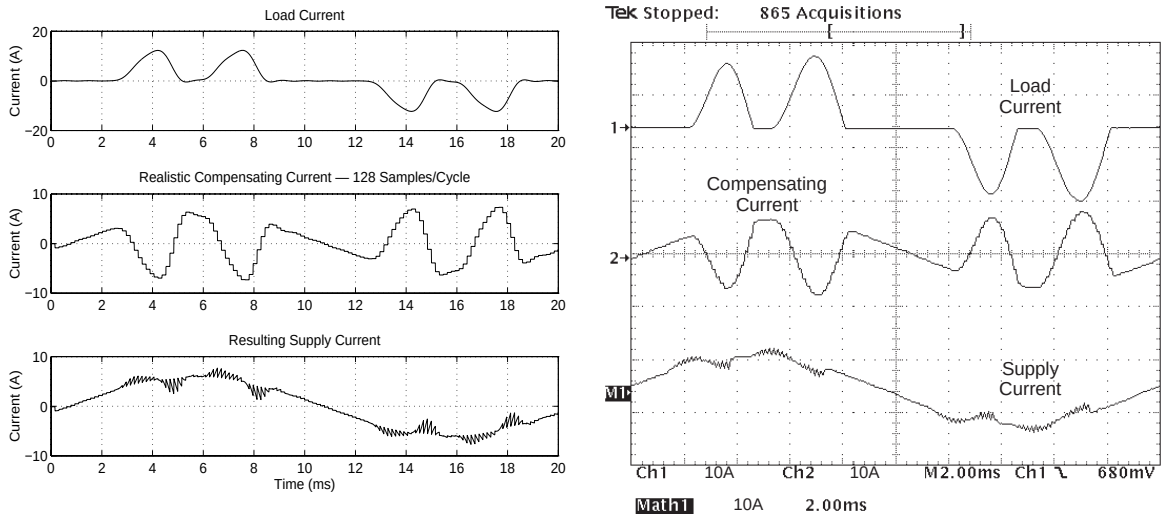

(a) (b)

Figure 8.6 Comparison of (a) the simulated effects of sampling and quantising to those (b) observed experimentally.

It can be seen from the results in Figure 8.6 that the simulation and experimental results are almost identical. This gives confidence that the Simulink model of the active filter controller is now accurate.

8.3.3 The Effect of Sampling Frequency

Further simulations were made with Simulink to ascertain the effect of sampling frequency on the delay artefacts. The same model was used for each simulation and the only parameter changed was the number of samples per cycle of mains current. The number of samples used were 128 (as for the DSP implementation), 256, 512 and 4096 (to approximate a continuous system). Figure 8.7(a)–(d) show the effects of different sampling frequencies on resulting supply currents.

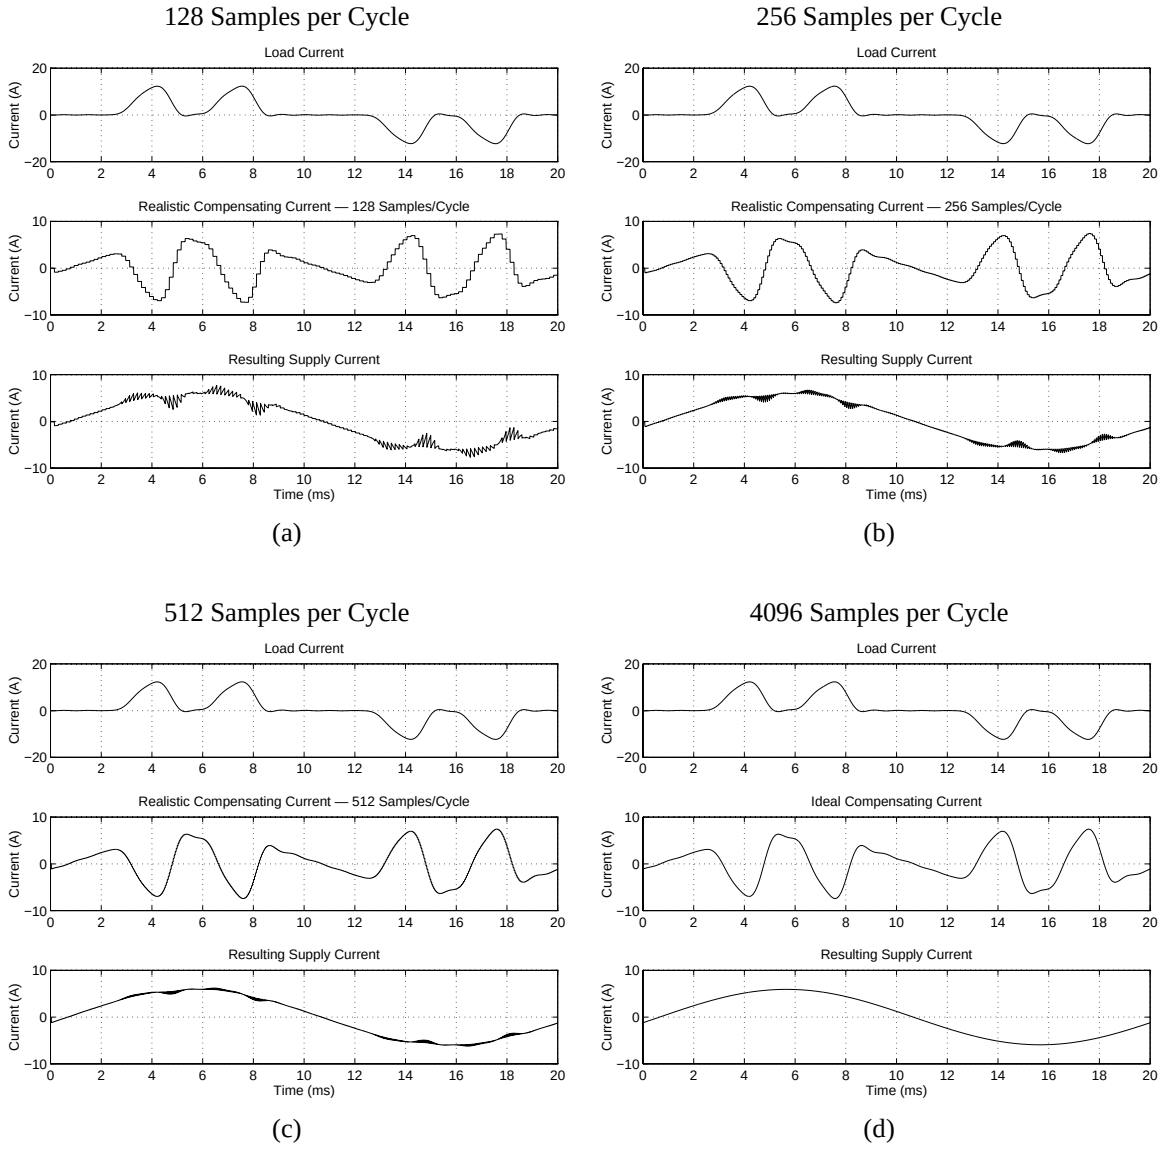

Figure 8.7 Simulation results for: (a) 128 samples, (b) 256 samples, (c) 512 samples and (d) 4096 samples.

The supply current became more sinusoidal as the sampling frequency was increased. The supply current waveform in Figure 8.7(c) is almost a perfect sinusoid and is a marked improvement over the supply current in Figure 8.7(a). When the sampling is fast enough to be considered continuous, the supply current is sinusoidal, as seen in Figure 8.7(d). A sampling rate of 512 samples per mains cycle has been used in previous single phase active filter work (Round, 1992) but would only give 217 cycles per sample (instead of the present 868 cycles per sample) for a three phase active filter controller. The 256 tap FIR high pass filter and sinusoidal subtraction methods would not operate at this sampling frequency as the processor load of 122% exceeds the capability of the DSP. The other techniques would have processor

loads ranging from 31% for notch filtering to 62% for the 128 tap high pass filter. A sampling frequency of 256 samples per cycle would be a suitable compromise between high processor loadings and the effects of the sample hold times. The processor loads are manageable and the ‘notching’ is markedly reduced in Figure 8.7(b) for this sampling frequency.

8.4 Quality of Compensation

The steady state effectiveness of an active filter can be determined by looking at the Total Harmonic Distortion (THD) of the load current and the THD of the resulting supply current. Ideally the THD of the supply current will be 0%, indicating that no harmonic currents are present. It is important that no undesired 50Hz components are present in the compensating current as this represents an injection of real power. The flow of real power must be carefully controlled to maintain the active filter inverter’s DC bus at a near constant voltage. The quality of compensation for each method was examined by looking at 50Hz rejection and the reduction in THD of the supply current with three non-linear loads.

8.4.1 Rejection of 50Hz Currents

A Hewlett-Packard Dynamic Signal Analyser (DSA) measured the magnitude of the fundamental component of the load and compensating currents. The difference between the two is the 50Hz rejection and Table 8.2 summarises the results of this test. The method that provided the greatest rejection of the fundamental was the FFT based technique, closely followed by the 256 tap High Pass FIR filter. Even so, the best result was only a 99% reduction in the level of the fundamental.

Table 8.2 50Hz rejection capability of each harmonic isolation method

	50Hz Input (dBV)	**50Hz Output** (dBV)	**Attenuation** (dB)	**Leakage** (%)
128 Tap FIR High Pass	-19.45	-43.06	23.6	6.6
256 Tap FIR High Pass	-19.46	-58.35	38.9	1.1
IIR Notch	-19.45	-52.55	33.1	2.2
Sinusoidal Subtraction	-19.47	-41.75	22.3	7.7
Fast Fourier Transform	-19.45	-59.08	39.6	1.0

8.4.2 Three Phase Bridge Rectifier

A three phase uncontrolled diode bridge rectifier generated non-linear load currents. These currents were fed into the active filter controller with the load current measurement board described in §4.4. A variable resistance load bank was connected to the DC side of the rectifier to vary the current drawn from the AC supply. The currents were altered for each test by changing the load on the DC side of the rectifier. The waveform captures shown in Figure 8.8–Figure 8.10 are from the PC software used to monitor the active filter's performance.

8.4.2.1 No Filtering

For the first test, the load bank was connected directly to the DC terminals of the rectifier and Figure 8.8 shows the waveform of the current drawn from the AC supply.

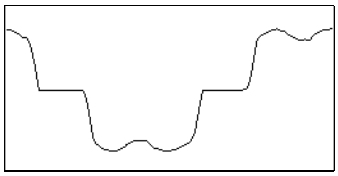

Figure 8.8 Current waveshape with a purely resistive DC load.

Six load settings were used on the load bank: "1A", "2A", "4A", "8A", "16A" and "20A". These were the nominal RMS currents when the load bank was connected to a 240V mains supply and bore no relevance to the DC current drawn from the rectifier. Table 8.3 summarises the results of this testing.

Table 8.3 Effectiveness of compensating an unfiltered three phase bridge rectifier.

Load Current						
Load Bank Setting	"1A"	"2A"	"4A"	"8A"	"16A"	"20A"
AC Load Current (I_L)	0.29A	0.58A	1.14A	2.26A	4.39A	5.43A
AC Load Current THD	33.0%	30.3%	29.6%	29.1%	27.9%	27.5%
Supply Current THD						
128 Tap FIR High Pass	12.0%	5.3%	2.8%	2.2%	1.2%	0.8%
256 Tap FIR High Pass	8.4%	4.3%	2.1%	1.9%	1.0%	0.8%
IIR Notch	1.1%	0.7%	0.6%	0.5%	0.5%	0.4%
Sinusoidal Subtraction	7.2%	3.1%	1.5%	1.2%	0.7%	0.5%
Fast Fourier Transform	9.6%	4.8%	2.2%	1.9%	1.1%	0.8%

The THD of the supply current is reduced as the load current increases. This is due to the sampled load current using ever increasing amounts of the analogue to digital converter's dynamic range, reducing the effects of noise.

8.4.2.2 Rectifier with Capacitive Voltage Smoothing

The second test involved placing 10000μF of capacitance across the DC output of the three phase rectifier. This smoothed the DC output voltage but caused the rectifier to draw current from the AC side in pulses. This type of configuration is common in the front-end power supply of variable speed motor drives (Bose, 1997). The high current peaks are shown in the waveform capture in Figure 8.9 and are due to large current inrushes into the capacitors once the rectifier output voltage exceeds that on the capacitors.

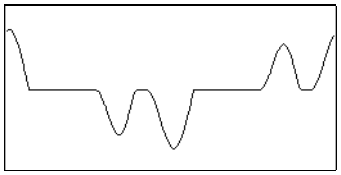

Figure 8.9 AC load current with a capacitive DC load.

The RMS currents used for testing the active filter controller were reduced because of the increased crest factor of the waveform. The active filter controller would not generate a correct compensating current if the ADC input saturated. Four load settings were used on the load bank and Table 8.4 summarises the testing results.

Table 8.4 Effectiveness of compensating for a three phase bridge rectifier with capacitive smoothing.

Load Current				
Load Bank Setting	"1A"	"2A"	"4A"	"8A"
AC Load Current (I_L)	0.53A	0.94A	1.72A	3.08A
AC Load Current THD	156.2%	133.7%	112.6%	90.7%
Supply Current THD				
128 Tap FIR High Pass	13.4%	10.0%	5.4%	2.6%
256 Tap FIR High Pass	12.3%	5.8%	3.3%	1.7%
IIR Notch	2.6%	2.1%	1.8%	1.5%
Sinusoidal Subtraction	7.3%	4.5%	2.2%	1.1%
Fast Fourier Transform	11.7%	6.4%	4.1%	2.1%

The results are similar to the previous test, with the THD of the resulting supply current decreasing as larger load currents are used. The quality of compensation is not as good due to the delay artefacts discussed in §8.3. The load current peaks increase the slope of the current, making the delay effects more pronounced.

8.4.2.3 Rectifier with Inductive Current Smoothing

The current drawn from the DC output of a three phase rectifier was smoothed by placing a 120mH inductor in series with the load. The smoothed DC current gave flat currents on the AC side. Extremely inductive loads are often presented by large DC motors used for traction and similar applications. Figure 8.10 shows the inductive current waveshape that results from the current smoothing inductor on the DC side of the rectifier.

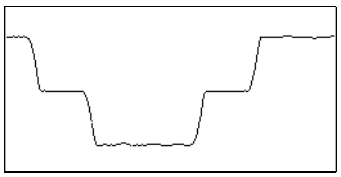

Figure 8.10 Inductive current smoothing on the DC side gives a flat AC side current.

Six settings were used on the load bank and Table 8.5 summarises the results of this testing. As with the two previous tests, a larger load current resulted in a lower supply current THD. The THD of the output is lower than for the capacitively smoothed test, but is very similar to the test with no filtering.

Table 8.5 Effectiveness of filtering a three phase bridge rectifier with inductive current smoothing.

Load Current						
Load Bank Setting	"1A"	"2A"	"4A"	"8A"	"16A"	"20A"
AC Load Current (I_L)	0.30A	0.59A	1.13A	2.25A	4.29A	5.27A
AC Load Current THD	32.7%	30.3%	29.3%	28.2%	26.7%	26.0%
Supply Current THD						
128 Tap FIR High Pass	13.5%	7.5%	3.1%	2.0%	1.4%	0.9%
256 Tap FIR High Pass	10.8%	5.1%	2.3%	1.9%	1.2%	0.8%
IIR Notch	1.1%	0.7%	0.6%	0.4%	0.4%	0.5%
Sinusoidal Subtraction	6.6%	3.8%	1.9%	1.2%	0.7%	0.6%
Fast Fourier Transform	11.9%	6.5%	2.6%	1.9%	1.3%	0.8%

8.4.3 Quality of Compensation Summary

The steady-state results from the active filtering tests performed here show that Notch Filtering gives the best harmonic compensation across all current ranges, with capacitive, inductive and no filtering of the rectifier DC output. The ranked order of compensation quality is:

1. Notch Filtering (IIR)
2. Sinusoidal Subtraction (256 tap FIR)
3. High Pass Filtering (256 tap FIR)
4. Fast Fourier Transform
5. High Pass Filtering (128 tap FIR)

The differences among the techniques are most noticeable at low currents. This may be due to the reduced signal to noise ratio from the relative increase in quantisation noise. For the case of inductively smoothed DC currents, the notch filter method reduced the THD by 31.6% at $0.30A_{RMS}$ and 25.5% at $5.27A_{RMS}$. The reduction in THD for the 256 tap High Pass Filter was more extreme; the THD was reduced by 21.9% at $0.30A_{RMS}$ and 25.2% at $5.27A_{RMS.}$ Programmable gain transimpedance amplifiers may reduce the noise present at the ADC input and improve the quality of compensation for low current ranges.

8.5 Transient Performance

Simulation results in Chapter 3 showed that each harmonic isolation method reacted differently to changes in load current. Two tests were performed to assess the ability of each method to deal with transients. Firstly, a single step change was applied to the load current by increasing the load on the DC side of the three phase bridge rectifier. The resulting supply current was then allowed to stabilise and the waveform recorded over this period. Secondly, the load current was rapidly changed to mimic the regular changes that would be experienced in an industrial or commercial environment.

8.5.1 Response to Step Changes in Load

The load current for the active filter was increased by changing the resistive load presented to the three phase bridge rectifier with no additional smoothing. The load bank was initially set to “1A” and then switched to “10A” + “1A” to give a significant step change. A timebase of 20ms per division was used for all waveform captures, except for IIR Notch Filtering where 50ms per division was required because of its slow response. The results of the five tests are shown in Figure 8.11.

The step change test shows that each harmonic isolation method behaves differently. The two FIR high pass filtering techniques have very fast transient responses, with the 128 tap filter responding in one fundamental cycle and the 256 tap filter responding in two cycles.

The IIR notch filter mimics an analogue Butterworth filter and has the same slow transient response. Ten cycles elapsed before the magnitude of the resulting supply current stabilised. During this time the compensating current would have had a large fundamental component that in turn would have collapsed the bus voltage on the active filter inverter's DC bus. This increases the RMS current injected by the active filter and so an inverter with a larger VA rating will be required to deal with such transient load changes.

Sinusoidal subtraction had a similar transient response to the 256 tap FIR high pass filter. This is expected because a 256 tap FIR low pass filter was used to recover the 50Hz component of the current. The compensating current is incorrect until the magnitude of the fundamental is updated at the next half cycle peak. The transient response is no better than that of the 256 tap FIR high pass filtering method.

Mathematical assumptions made by the Fast Fourier Transform method mean that the supply current is distorted immediately after a change in load. The compensating current was correct once the FFT routine had a chance to perform one of the cycles shown in Figure 8.2.

8.5.2 Rapid Load Changes

The single step change shows what would happen when a piece of machinery was powered up, but not what would happen if the load current was constantly changing. To obtain some indication on the effect of rapid changes in load, the resistive load bank was switched so the current was constant for only a few cycles of 50Hz mains. Figure 8.12(a)–(e) shows the response for each harmonic isolation method to rapidly changing load currents.

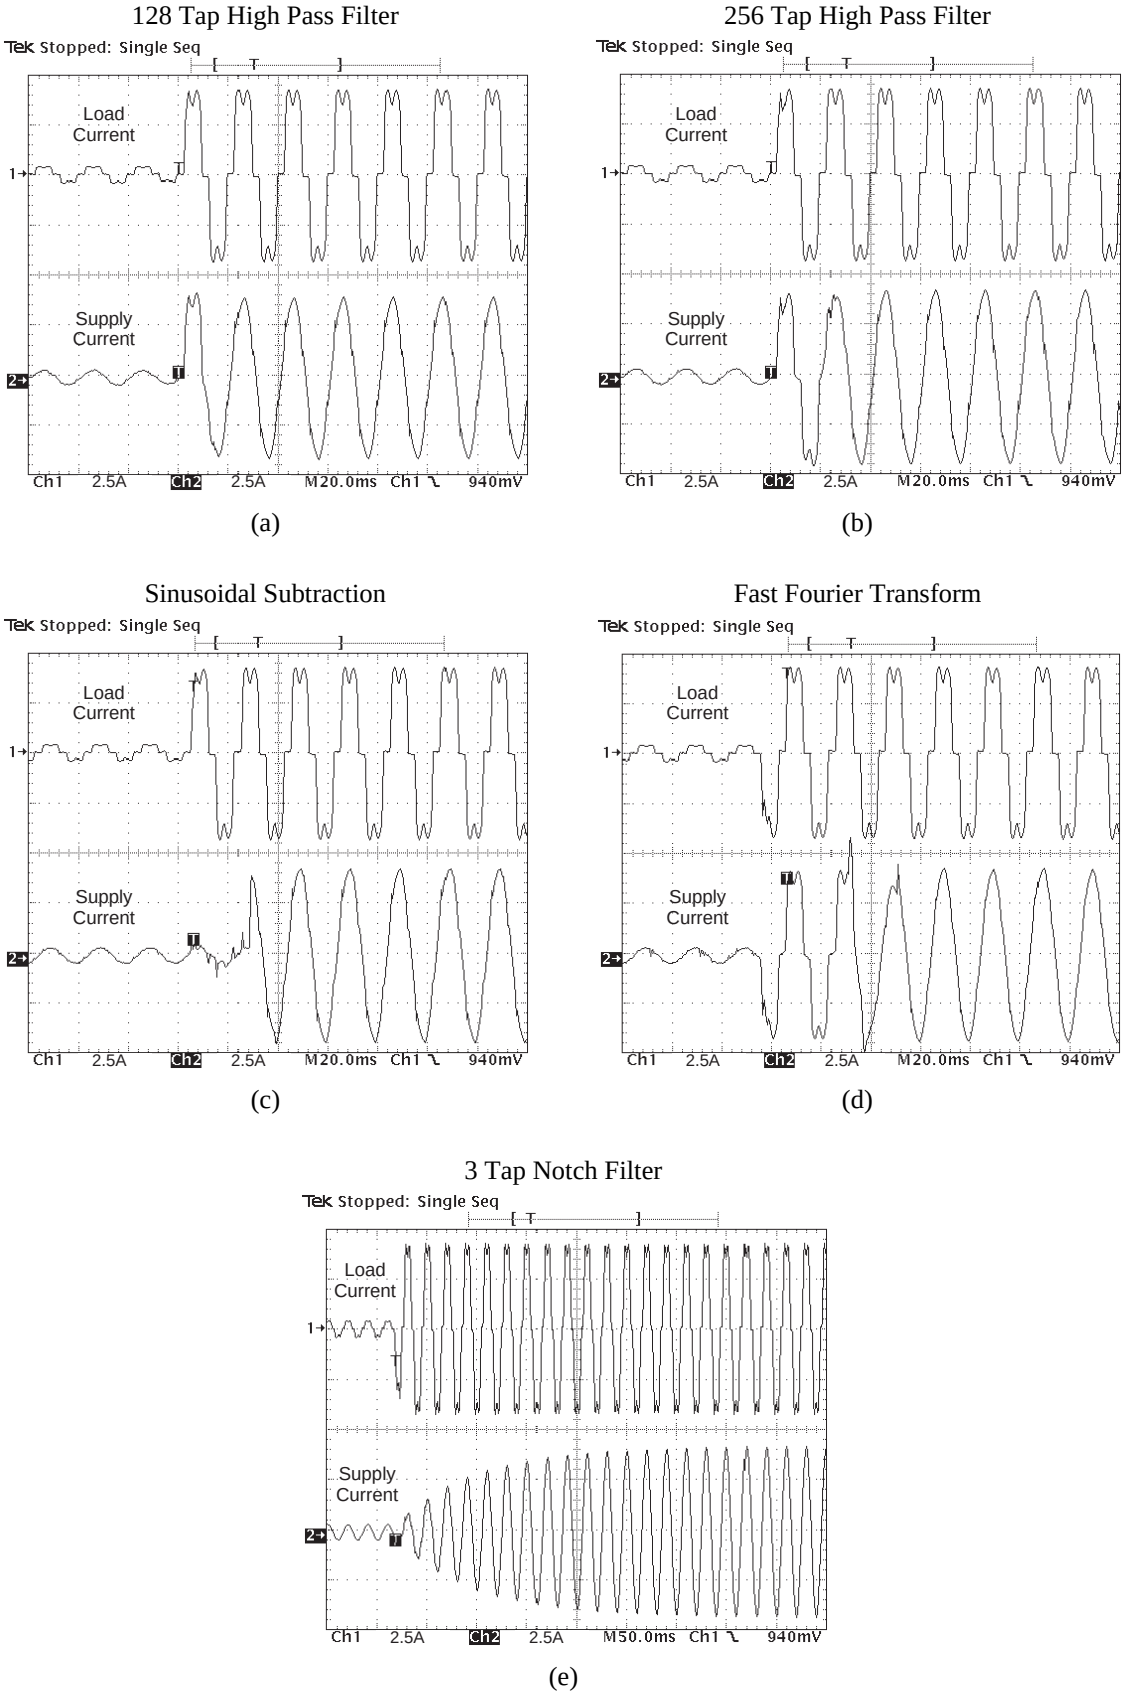

Figure 8.11 Transient response to a step change in load for: (a) 128 tap FIR high pass filter, (b) 256 tap FIR high pass filter, (c) Sinusoidal Subtraction, (d) Fast Fourier Transform and (e) IIR notch filter.

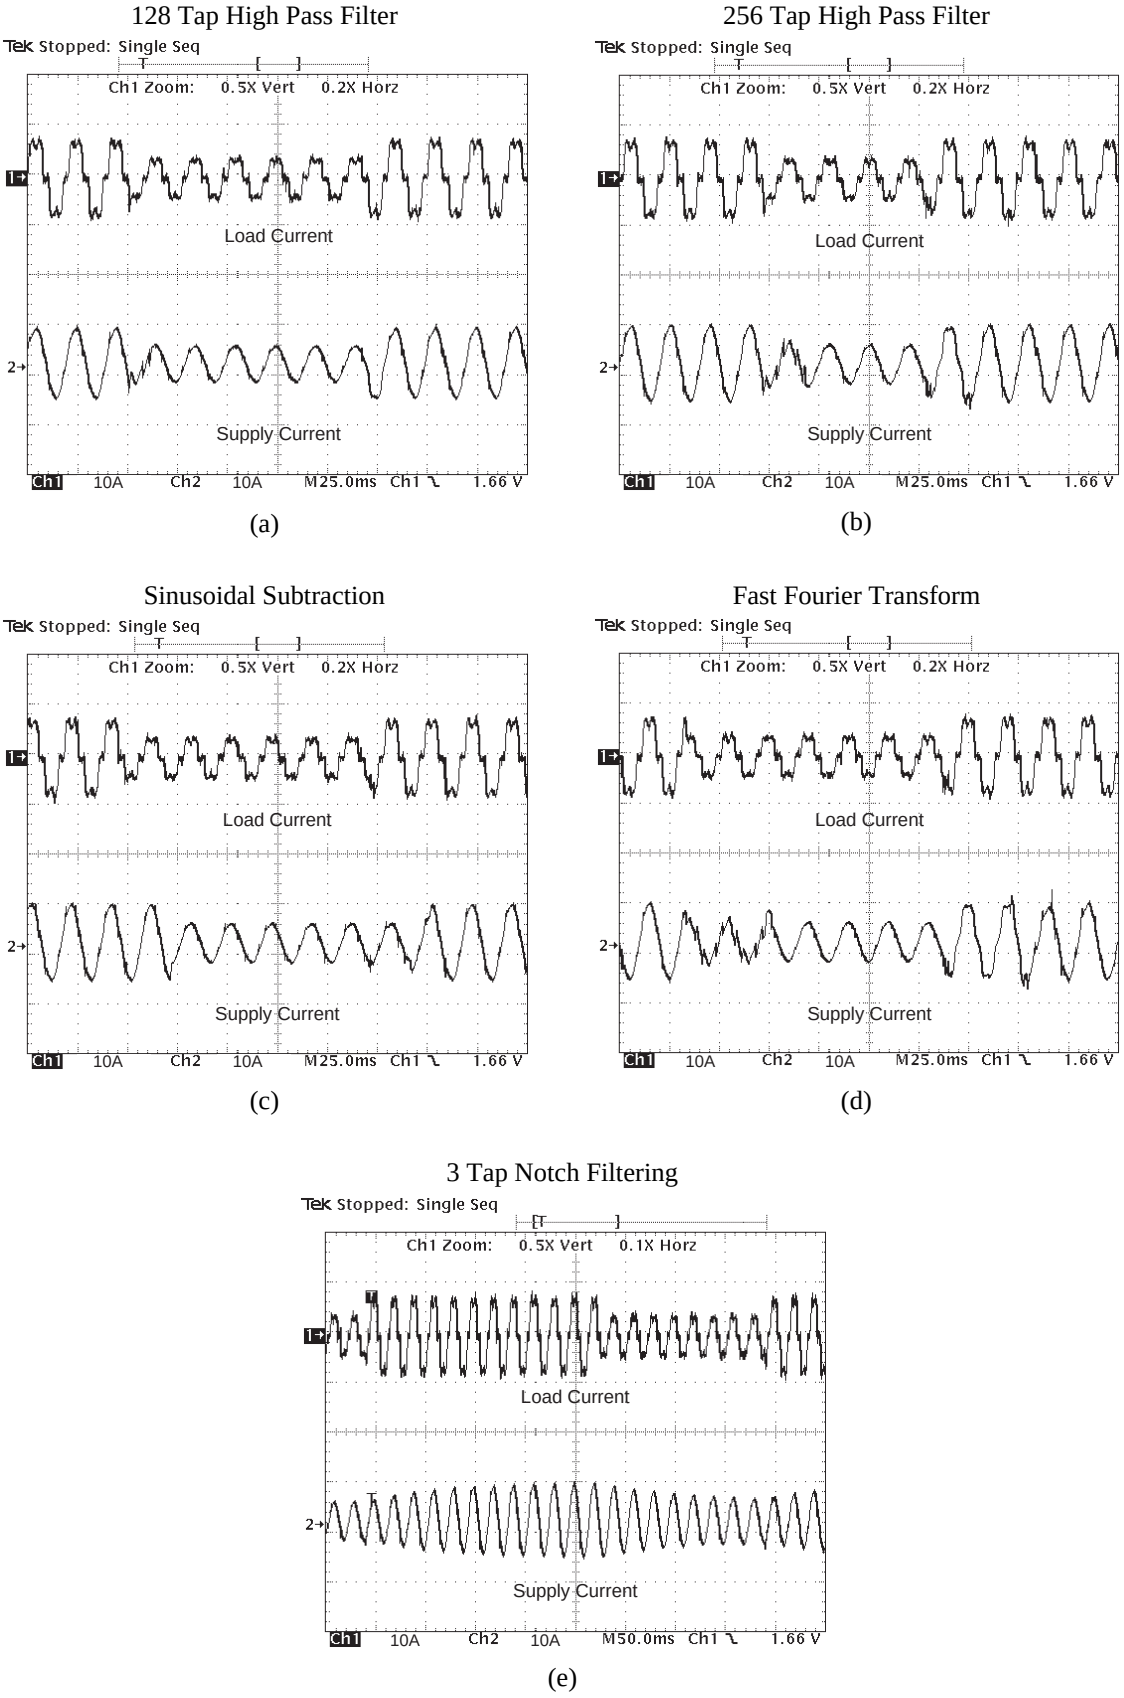

Figure 8.12 Response to rapid changes in load for: (a) 128 tap FIR high pass filter, (b) 256 tap FIR high pass filter, (c) Sinusoidal Subtraction, (d) Fast Fourier Transform and (e) IIR notch filter.

The two methods that responded quickest to a single step change in load, the 128 tap and 256 tap FIR high pass filters, again showed fast response times. Sinusoidal subtraction gave the cleanest result because the waveshape of the sinewave being subtracted was not altered by the step changes. The IIR notch filter was unable to cope with the rapid load changes due to its slow transient response. Rapid changes in load current cause the Fast Fourier Transform method to give distorted supply currents because two mains cycles are required for its calculations, reducing the effectiveness of this method.

8.6 Selective Harmonic Cancellation

The Fast Fourier Transform (FFT) based harmonic isolation method offers interesting alternatives to total harmonic elimination.

8.6.1 Concept

Recent publications have looked at the connection of more than one active filter to a single bus (Morán *et al.*, 1993; Morán *et al.*, 1997). These methods have generally involved one active filter compensating as much as possible, with another active filter completing the compensation. A disadvantage of this technique is that the division of compensating current between the two filters is not defined. The technique proposed in this section overcomes this and offers more flexible control of multiple active filters.

The FFT method eliminates the fundamental component from the load current by setting the real and imaginary parts of the 50Hz component to zero. The IFFT then recreates the time domain waveform without the 50Hz component. This technique can be applied to other harmonics, preventing their elimination. It is quite possible for the FFT method to compensate for the fifth, seventh and ninth harmonics and leave all others in the resulting supply current.

8.6.2 Uses of Selective Cancellation

8.6.2.1 Multiple Active Filters

Selective harmonic cancellation lends itself well to installations with more than one active filter. Figure 8.13 shows in schematic form an installation with two active filters. In this example Active Filter One compensates from the second harmonic up to the tenth and Active Filter Two takes over and compensates from the eleventh to the fiftieth harmonic.

The bulk of the compensating current will be in the lower order harmonics and will be generated by Active Filter One. The relatively low upper limit of compensation of 500Hz for this active filter allows the inverter switching frequency to be reduced. This reduces losses in the semiconductor switches and improves the overall efficiency. Active Filter Two compensates for higher order harmonics with relatively small magnitudes. The reduced compensating current means smaller semiconductors, optimised for higher switching frequencies, can be used. This leads to increases in efficiency due to lower switching losses.

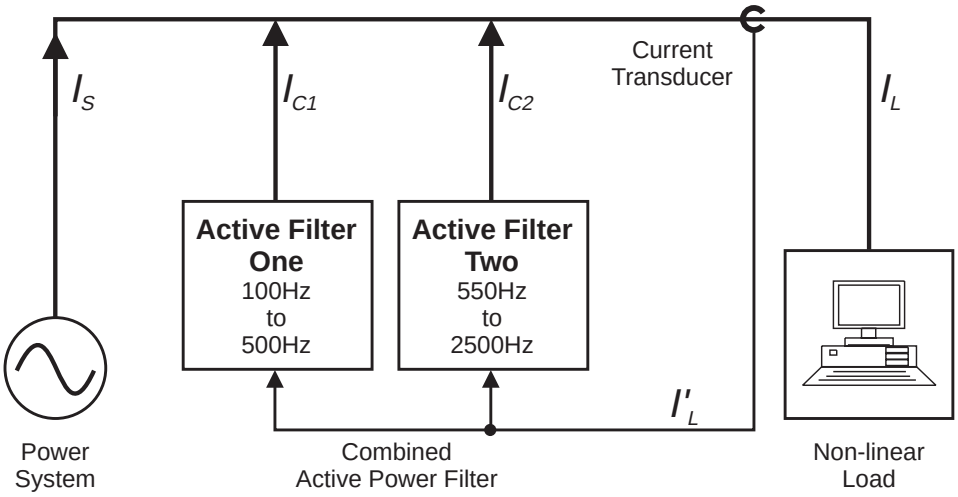

Figure 8.13 Installation with two active filters. Active Filter One is compensating the low harmonics (2^{nd} to 10^{th}) and Active Filter Two is compensating for the higher order harmonics (11^{th} to 50^{th}).

Simulations performed in Simulink show the advantages of using two active filters for harmonic compensation. Figure 8.14(a) shows the compensating current generated from a single active filter with a hysteresis current controlled output. The switching frequency is high and leads to unnecessary switching on the low current gradients. Figure 8.14(b) is the sum of the compensating currents generated by two active filters. The low frequency components (from 100Hz to 350Hz) are shown in Figure 8.14(c) and the high frequency components (from 400Hz to 1000Hz) shown Figure 8.14(d). The inductance used for the single active filter and the high frequency active filter was 700μH. The output current in Figure 8.14(c) was limited to 350Hz and so an injection inductance of 1400μH was used. This lowered the average switching frequency for the bulk of the current from 16kHz in Figure 8.14(a) to 11kHz in Figure 8.14(c).

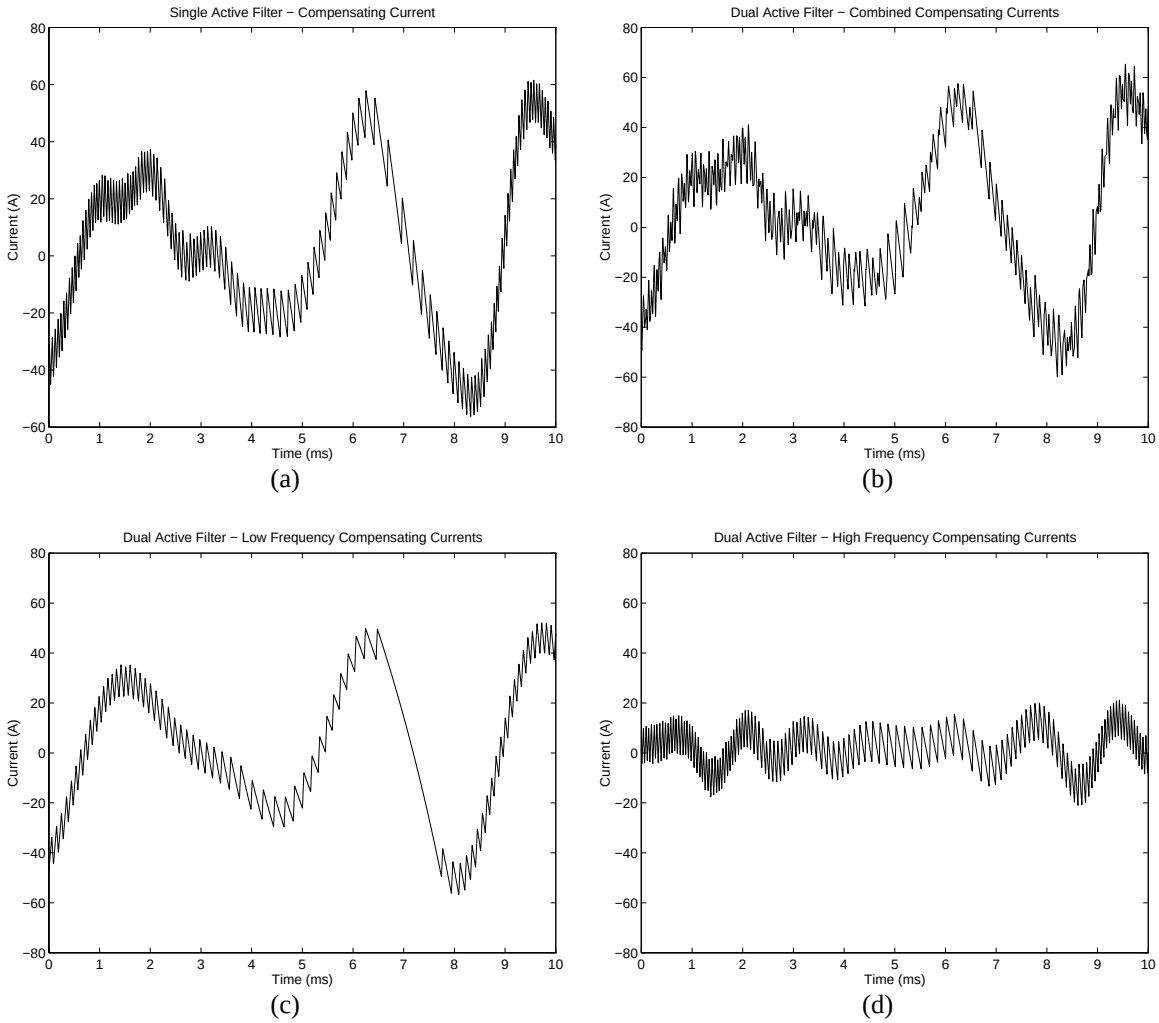

Figure 8.14 Simulation of a single active filter system (a) and a dual active filter system (b). Plots (c) and (d) are the low frequency and high frequency components of (b).

8.6.2.2 Self Limiting Active Filters

All active filters have a maximum power rating determined by the inverter's thermal characteristics. If the demand for harmonic compensation is likely to exceed the maximum rating of the inverter, the compensating current must be reduced to prevent damage to the power semiconductors, or the active filter shutdown.

A simple way of limiting the VA output of an active filter is to scale the entire compensating current until the RMS current is the desired value. An alternative provided by selective harmonic isolation is to limit the number of harmonics that are compensated. Large low frequency harmonics can be completely eliminated, leaving higher order harmonics in the supply current. If on the other hand, higher order harmonics cause the most trouble, such as in

telecommunication systems, these would be compensated first and the lower order harmonics left uncompensated. Simple scaling is still possible with the FFT harmonic isolation method and the two techniques can be combined through a set of adaptive rules to give the best compensation.

8.6.3 Experimental Examples of Selective Cancellation

The FFT harmonic isolation method implemented in the DSP was modified to selectively cancel harmonic currents. This example shows how selective cancellation can be used with two active filters. One active filter cancels all harmonics up to and including the eighth and another cancels harmonics from the ninth to the fiftieth. The PC software screen captures of Figure 8.15 and Figure 8.16 show the separation of the harmonic compensation into two bands. The PC displayed the load current and supply current as a time domain waveform and as a frequency spectrum. For this example, comparison of the load and supply current spectra gives the most information.

The upper left plot in each figure is the time domain representation of the currents. The light grey waveform is the load current seen by the active filter. The true resulting supply current is shown in mid grey and the calculated supply current in black. The upper right plot shows the frequency spectrum of the load current. The two lower plots are the frequency spectra of the supply current, the left plot being the spectrum of the true current and the right plot that of the calculated supply current. The display in the middle of the screen indicates the Total Harmonic Distortion (THD) and RMS current for the load current and two supply currents.

8.6.3.1 Cancellation of Harmonics up to and Including the Eighth

The majority of harmonic currents generated by a six pulse rectifier will be the fifth and seventh harmonics. The eighth harmonic, 400Hz, was arbitrarily selected as a cut-off point between the high magnitude, low frequency harmonics, and the low magnitude, high frequency harmonics. A screen capture of the PC monitoring software (described in §6.4.3) is shown in Figure 8.15.

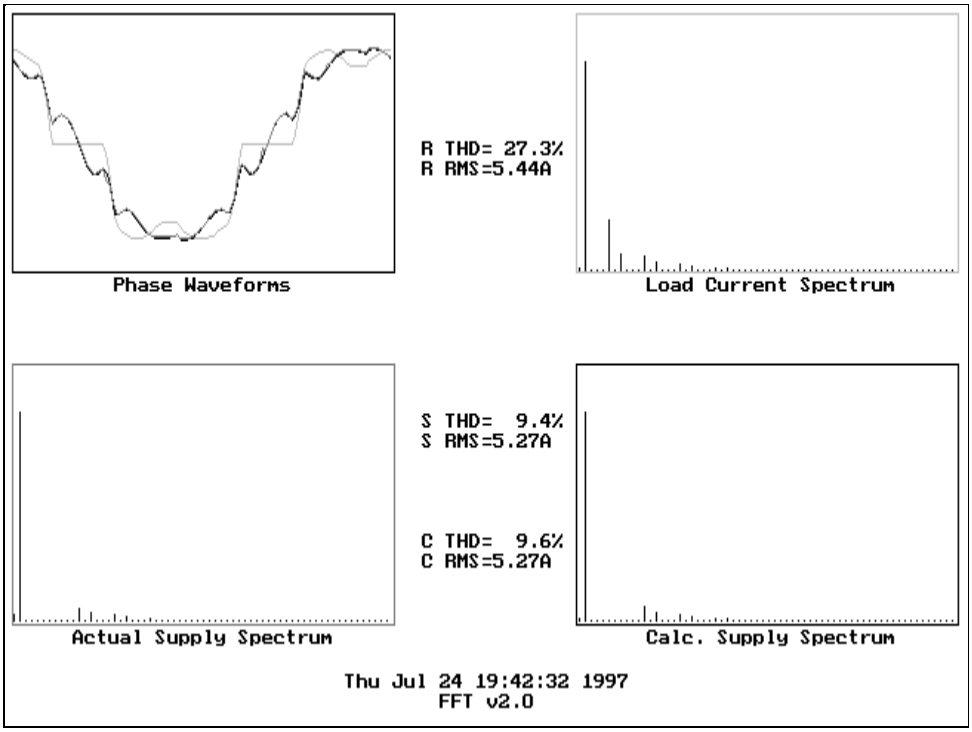

Figure 8.15 The results of compensating for harmonic currents up to and including the eighth.

The supply current resulting from this cancellation is more sinusoidal than the load current and consequently has a lower THD. The supply current spectra show that the lower order harmonics have been removed, with the higher order harmonics remaining in the supply current.

8.6.3.2 Cancellation of Harmonics from the Ninth Upwards

The load current spectrum in Figure 8.15 shows that the majority of harmonics from a six pulse rectifier are below the eighth, though there are some higher harmonics that distort the supply current. This second example of selective harmonic cancellation continues from where the cancellation of the previous section stopped. The selective cancellation parameters were changed so that harmonics from the ninth to the fiftieth were cancelled, leaving the lower order harmonics in the supply current. A screen capture from the monitoring software is shown in Figure 8.16.

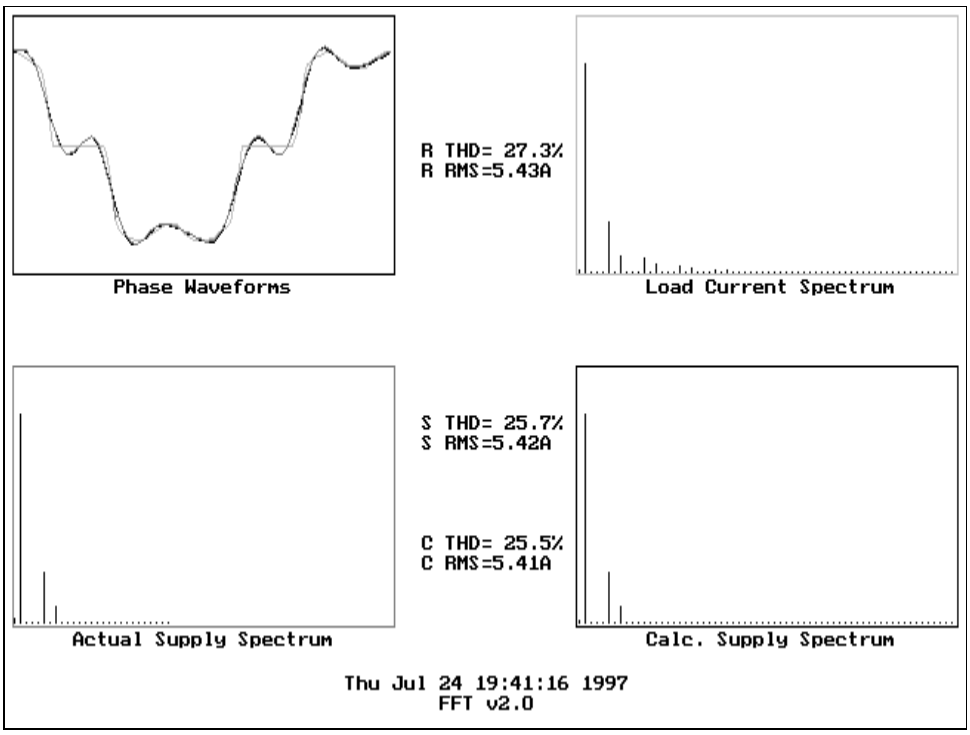

Figure 8.16 The result of compensating only for harmonic currents above the eighth.

The four plots in Figure 8.16 show the same type of harmonic current information as Figure 8.15. The supply current is far less sinusoidal and is reflected in the increased value of the supply current THD (26%, compared to 10% for Figure 8.15). The two supply current spectra show the presence of fifth and seventh harmonics. These are below the ninth harmonic and so no compensation has been performed.

8.7 Summary

This chapter has presented the results of an experimental evaluation of four different harmonic isolation methods, with the High Pass Filtering technique having two digital filter variants. An oversight in the simulation model used in Chapter 3 was discovered and corrected. The simulation and experimental results are now consistent, as shown by Figure 8.6, and this provides confidence in the simulation model.

The harmonic isolation methods have been evaluated by examining the computational load of the DSP, the quality of steady state harmonic compensation and the response to changes in load current. An analogue model of an active filter inverter was used to simulate current injection into a power system. The results of this testing summarised in Table 8.6. The processor load of each method was measured by counting the number of DSP instruction

cycles required to perform the calculations. Table 8.6 shows that the processor load varies greatly between the different techniques, from 7.9% for a single biquad IIR Notch Filter to 30.5% for a 256 tap FIR High Pass Filter.

Table 8.6 Summary of experimental results for five harmonic isolation techniques.

Method	**Processor Load**	**Fundamental Leakage**	**Steady State** (Ranking)	**Transient Response** (Ranking)	**Rapid Load Changes** (Ranking)
128 Tap HPF	15.5%	6.6%	5	1	1
256 Tap HPF	30.5%	1.1%	3	2	2
IIR Notch	7.9%	2.2%	1	5	5
Sine Sub	30.5%	7.7%	2	3	3
FFT	10.1%	1.0%	4	4	4

The Notch Filtering method gave good steady state performance and had a low processor usage. Unfortunately the IIR notch filter had a slow transient response and therefore did not respond well to changes in load current. A larger rating inverter in the active filter would be required to inject the fundamental current that would result from the incorrect compensating current signals. The 256 tap high pass filter gave reasonable steady state performance and excellent transient performance, at the expense of a high processor loading. The Fast Fourier Transform does not give the best results in either steady state or transient tests, but is consistently good in both. Additional functionality provided by this technique allows more sophisticated harmonic compensation to be performed without sacrificing its low computational complexity. Selective harmonic cancellation has been demonstrated through simulation and experimentation. This novel technique allows the connection of multiple shunt active filters in parallel without any interaction of injection currents.

Chapter 9
Discussion

9.1 Comparison of Simulation and Experimental Results

The simulations conducted as part of this work concentrated on the ability of several harmonic isolation methods to determine the harmonic content of unbalanced three phase load currents correctly. Such load currents are expected in industrial and commercial premises, the target market for the active filter under development (Emanuel *et al.*, 1993). Simulations of the steady quality of compensation are consistent with observed results. The Fast Fourier Transform (FFT) technique provides the most flexibility with its ability to treat harmonics individually, rather than as a group. This selective cancellation allows an active filter to compensate only those harmonics causing interference or damage to electrical equipment, reducing the capital cost of the active filter. The reliable connection of multiple shunt active filters in parallel is made possible by selective harmonic cancellation.

Initially an oversight was made in the design of the Simulink system model used for evaluating harmonic isolation techniques. The consequence was that the simulation results were better than those observed experimentally because delays in the digital controller were not taken into account. The model was then corrected and shown to produce the expected results. Increasing the sampling rate of the load current would reduce the effect of such delays, but at the penalty of increased processor loading. A compromise must be made between the quality of compensation required and the cost and complexity of the active filter controller.

The experimental results showed that those obtained by simulation were accurate. An IIR filter's transient response will be slow in comparison to the other techniques due to its recursive nature. The transient response of an FIR filter is dictated by the length of the filter, with the 128 and 256 tap filters having the expected response times of one and two mains cycles respectively. The FFT method has a transient response time of three cycles because of its cyclic calculations (discussed in detail in §8.2.4).

The agreement between the simulated and experimental results removes the need to test every harmonic isolation method experimentally. Simulation results from Simulink have been validated for a variety of harmonic isolation methods. The results from simulating the

digital hysteresis current controller match closely those observed experimentally. This gives added confidence in the ability to accurately model and simulate an entire active filter, from load current measurements through to the injection of a compensating current. Simulation gives the designer much finer control of the operation of the system under test and the ability to investigate certain behaviours more closely than would be possible with an experimental active filter. A digital control system presents many advantages over an analogue active filter controller, but at the expense of cost and complexity. The optimisation of the operating parameters for a digital controller is an involved task, much more suited to simulation than experimentation.

9.2 Future Work

9.2.1 Operation of a Digitally Controlled Active Filter

The current work has involved the experimental evaluation of an active filter controller and an inverter controller. It was intended for these to be connected to produce a completely digital active filter, but insufficient time was available to complete this. Once the interfacing of the active filter controller to the inverter is proven, an embedded controller using the low-cost TMS320C31 DSP could be built. This would be used in a standalone active filter without a PC. This embedded controller would be based around the ’C31 DSP and a Xilinx FPGA, providing control for the entire active filter.

9.2.2 Supply Side Active Filtering

All of the harmonic isolation methods examined in this thesis have been based on load current measurements. Supply side active filtering uses current measurements ‘up stream’ from the inverter to control the injection of compensating currents. This connection allows the active filter controller to use existing metering current transformers (CTs) for harmonic current measurements. Using existing CTs reduces disruption to existing power connections and lowers the capital and installation costs of an active filter.

Supply side active filtering control methods are generally closed loop systems and may therefore suffer from instability. The simulation models developed as part of this work will allow reliable simulations of supply side active filter operation to be performed before a filter is commissioned at a new installation.

9.2.3 Multilevel Inverter Controller

The present digital hysteresis current controller is for a two level inverter, with the output of the inverter swinging between the positive and negative DC supplies. A multilevel inverter has three or more output voltages and reduces the voltage stresses on the individual semiconductors (Choi, Cho and Cho, 1991). The three level inverter, also called a Neutral Point Clamped inverter, offers a zero output voltage as well as positive and negative outputs. This can be used to reduce the current slope of the injected compensating current where appropriate, resulting in a lower average switching frequency. By reducing the number of switchings, the losses in the power inverter will be reduced and the active filter will become more efficient.

The decision making required for a multilevel current controlled inverter can be implemented in an FPGA. The interface to this more advanced controller should be the same as the present two level controller, allowing direct replacement in an active filter when the development is complete.

9.2.4 Selective Harmonic Cancellation

The investigations of Chapter 8 into selective harmonic cancellation with the FFT method are only the beginning of what may be possible. Many novel techniques and control strategies can be based on the ability of the FFT harmonic isolation method to deal with each harmonic independently of all others. It was mentioned earlier in this chapter that an active filter may target specific harmonics, reducing the required inverter VA rating. Considerable cost savings can be made without seriously affecting the perceived performance of the active filter. The connection of many active filters in parallel permits distributed active filtering, with as many units as necessary coming ‘online’ to compensate harmonic currents. DSP based controllers will be able to communicate with one another, giving the best harmonic compensation for the lowest operating cost.

Chapter 10
Conclusion

The increased use of energy efficient electronics and advanced power electronics is adding to the harmonic pollution of the power system. It has been shown in the literature that shunt active filters are capable of removing these harmonic currents. Active filters have many advantages over traditional passive filters, especially in small low voltage installations. This thesis has presented a completely digital solution for harmonic compensation of three phase unbalanced loads.

A TMS320C30 Digital Signal Processor (DSP) was used to control the active filter. Synchronous sampling was used to prevent the Fast Fourier Transform (FFT) introducing errors due to spectral leakage. Synchronous sampling allowed the digital filter to track variations in the mains frequency without recalculating the filter coefficients. Three phase load current measurements were made with LEM active flux nulling current transducers. These had a flat frequency response from DC to above 2.5kHz, the maximum frequency of interest. The current outputs from the three LEM sensors were converted to voltages and sampled by a four channel 12 bit analogue to digital converter. A switching inverter was built to inject compensating currents into the power system. This used a Siemens IGBT module that contained six IGBTs and their anti-parallel diodes in a single package. The high thermal mass and robustness of this module made it particularly suitable for experimental inverter development.

The inverter was controlled by a novel digital hysteresis current controller implemented in a Xilinx XC4010 Field Programmable Gate Array (FPGA). This controller made all switching decisions digitally and interfaced directly to a high speed serial port on the active filter controller DSP. Digital inverter feedback was achieved with three LEM current sensors in the inverter output and three extremely high speed 12 bit serial analogue to digital converters that connected to the Xilinx current controller. The operation of this controller was simulated using Simulink to determine the optimum characteristics before designing the logic for the FPGA. High quality sinusoidal and harmonic currents have been produced by this inverter to show its suitability for active filtering applications. The digital interface between the DSP and the current controller FPGA is extremely resistant to external interference, and

this gives the digital current controller a significant advantage over traditional analogue controllers.

The software for the DSP was written in C, with time critical digital filtering routines hand-coded in assembly language. Two FIR high pass filters, an FIR low pass filter and an IIR notch filter were implemented. This gave a modular structure to the software, keeping the software for each experimental active filter controller as similar as possible. MATLAB's Signal Processing Toolbox was used to design the digital filters and to display the expected frequency response. A Dynamic Signal Analyser verified the experimental frequency response of the digital filters by swept sine analysis. This technique is often used to characterise analogue filters, but the work presented here has shown that digital filters may also be evaluated this way.

A review of load current harmonic isolation methods in the literature has been made. Notch Filtering, Sinusoidal Subtraction, Instantaneous Reactive Power Theory, Synchronous Reference Frame and Fast Fourier Transform methods were examined by simulation. Simulink was used to model the harmonic isolation techniques in an active filter controller. A non-linear load current was used to evaluate each harmonic isolation method. A change in harmonic content and magnitude was introduced to this load current. The supply current waveform and 'sliding THD' for each method were plotted to observe the transient response. Industrial and commercial sites generally have unbalanced load currents and so the operation of each method with an unbalanced load was examined. IRPT compensates correctly only when the neutral current is zero, and only occurs when the load is balanced and no triplen harmonics are present. The SRF method compensates correctly when triplen harmonics are present, but only if the load is balanced, otherwise it performs undesired load balancing. The FFT based technique dealt with all load currents correctly because no assumptions were made regarding the three phases. A novel technique was developed using the FFT and IFFT to cancel harmonic currents selectively. No other technique to date offers the ability to deal with each harmonic individually, rather than as part of a group.

An uncontrolled three phase diode bridge rectifier drew a non-linear load current from the mains supply. Capacitive and inductive filtering was added to the DC side of the rectifier to change the shape of the current waveform. This load current was used for the experimental evaluation of four harmonic isolation methods: high pass filtering, notch filtering, sinusoidal subtraction and the FFT. The computational complexity of each method was measured to gauge the processing power required by the active filter controller. The methods, in increasing

order of computational complexity are: Notch Filtering (7.9%), FFT (10.1%), 128 tap High Pass Filtering (15.5%), 256 tap High Pass Filtering (30.5%) and Sinusoidal Subtraction (30.5%). Steady state and transient measurements confirmed the results obtained through simulation, verifying the accuracy of the Simulink active filter controller model. Notch Filtering gave the best results for static loads, but the FFT method provided more complete compensation over a wide range of load current characteristics.

The three phase digital hysteresis current controller developed for this work performs the fastest sampling of injection currents known to date and is the only implementation to use an FPGA for all calculations. Accurate models of the active filter controller and digital current controller have been developed using Simulink to assist in these and future developments. IRPT and SRF are the most common harmonic isolation methods in the literature, but have been shown to be unsuitable for unbalanced three phase systems. Single phase techniques, in particular notch filtering and the FFT, provide accurate compensation of most load currents. Selective harmonic cancellation using the FFT harmonic isolation method is the most promising result arising from this work. Simulations and experimental work presented here have shown this technique to be a viable method of connecting multiple shunt active filters in parallel, with the potential to reduce the operating expense of active filters. This may be the development needed to take active filters out of the laboratory and into industry.

References

Akagi, H. (1992): "Trends in Active Power Line Conditioners". *18th International Conference on Industrial Electronics, Control and Instrumentation (IECON '92)*, San Diego, USA, pp. 19-24.

Akagi, H. (1996): "New Trends in Active Filters For Power Conditioning", *IEEE Transactions on Industry Applications*, Vol. 32, No. 6, pp. 1312-1322.

Akagi, H., Nabae, A. and Atoh, S. (1986): "Control Strategy of Active Power Filters Using Multiple Voltage-Source PWM Converters", *IEEE Transactions on Industry Applications*, Vol. IA-22, No. 3, pp. 460-465.

Aredes, M., Häfner, J. and Heumann, K. (1997): "Three-Phase Four-Wire Shunt Active Filter Control Strategies", *IEEE Transactions on Power Electronics*, Vol. 12, No. 2, pp. 311-318.

Aredes, M. and Watanabe, E.H. (1995): "New control algorithms for series and shunt three-phase four-wire active power filters", *IEEE Transactions on Power Delivery*, Vol. 10, No. 3, pp. 1649-56.

Arrillaga, J., Bradley, D.A. and Bodger, P.S. (1985): *Power System Harmonics*, John Wiley & Sons, Chichester.

AS 2279.1 (1991): *Disturbances in mains supply networks. Part 1: Limitation of harmonics caused by household and similar electrical appliances*, Standards Australia.

AS 2279.2 (1991): *Disturbances in mains supply networks. Part 2: Limitation of harmonics caused by industrial equipment*, Standards Australia.

Bernard, S. (1997): "Harmonic Pollution", *Power Quality Online*, http://www.powerquality.com/art0041/art1.htm, 25 May 1997.

Bhattacharya, S., Divan, D.M. and Banerjee, B. (1991): "Synchronous Frame Harmonic Isolator Using Active Series Filter". *EPE '91*, Firenze, Italy, Vol. 3, pp. 3.030-3.035.

Borland (1992): *Turbo C++ Version 3.0*, Borland International, Inc.

Bose, B.K. (1992): "Recent Advances in Power Electronics", *IEEE Transactions on Power Electronics*, Vol. 7, No. 1, pp. 2-16.

Bose, B.K., Ed. (1997). *Power Electronics and Variable Frequency Drives*. IEEE Press, New York.

Brigham, E.O. (1974): *The Fast Fourier Transform*, Prentice-Hall, Englewood Cliffs.

Brod, D.M. and Novotny, D.W. (1985): "Current Control of VSI-PWM Inverters", *IEEE Transactions on Industry Applications*, Vol. IA-21, No. 4, pp. 562-570.

BS EN 61000-3-2 (1995): *Electromagnetic compatibility (EMC) — Part 3: Limits — Section 2: Limits for harmonic current emissions (equipment input current less than or equal to 16A per phase)*, British Standards Institution.

Chambers, J.A., Tantaratana, S. and Bomar, B.W. (1996): "Digital Filters", in *The Electronics Handbook* (Whitaker, J. C. ed). CRC Press, Inc., Boca Raton, FL, USA, pp. 749-772.

Chassaing, R. (1992): *Digital Signal Processing with C and the TMS320C30*, John Wiley & Sons, New York, NY, USA.

Choe, G.-H. and Park, M.-H. (1988): "A New Injection Method for AC Harmonic Elimination by Active Power Filter", *IEEE Transactions on Industrial Electronics*, Vol. 35, No. 1, pp. 141-147.

Choi, N.S., Cho, J.G. and Cho, G.H. (1991): "A General Circuit Topology of Multilevel Inverter". *22nd Annual Power Electronics Specialists Conference (PESC '91)*, pp. 96-103.

Duffey, C.K. and Stratford, R.P. (1989): "Update of Harmonic Standard IEEE-519: IEEE Recommended Practices and Requirements for Harmonic Control in Electric Power Systems", *IEEE Transactions on Industry Applications*, Vol. 25, No. 6, pp. 1025-1024.

Duke, R.M. and Round, S.D. (1993): "The Steady-State Performance of a Controlled Current Active Filter", *IEEE Transactions on Power Electronics*, Vol. 8, No. 3, pp. 140-146.

Duke, R.M., Round, S.D. and Henderson, K.C. (1990): "An Active Filter for Current Distortion Compensation in Power Systems". *Fourth International Conference on Harmonics in Power Systems*, Budapest, Hungary, pp. 367–373.

Emanuel, A.E., Orr, J.A., Cyganski, D. and Gulachenski, E.M. (1993): "A Survey of Harmonic Voltages and Currents at the Customer's Bus", *IEEE Transactions on Power Delivery*, Vol. 8, No. 1, pp. 411-421.

Fujita, H. and Akagi, H. (1991): "A Practical Approach to Harmonic Compensation in Power Systems - Series Connection of Passive and Active Filters", *IEEE Transactions on Industry Applications*, Vol. 27, No. 6, pp. 1020-1025.

Fukuda, S. and Endoh, T. (1995): "Control Method for a Combined Active Filter System Employing a Current Source Converter and a High Pass Filter", *IEEE Transactions on Industry Applications*, Vol. 31, No. 3, pp. 590-597.

Grady, W.M., Samotyj, M.J. and Noyola, A.H. (1990): "Survey of Active Power Line Conditioning Methodologies", *IEEE Transactions on Power Delivery*, Vol. 5, No. 3, pp. 1536-1542.

Hayashi, Y. and Takahasi, K. (1991): "A Novel Control of a Current-Source Active Filter for ac Power System Harmonic Compensation", *IEEE Transactions on Industry Applications*, Vol. 27, No. 2, pp. 380-385.

Henderson, K.C. (1989): "The Development of a High Performance Active Filter for Current Distortion Compensation in Power Systems", M.E. Project Report, University of Canterbury, Christchurch, New Zealand.

Holtz, J. (1997): "Pulse Width Modulation for Electronic Power Conversion", in *Power Electronics and Variable Frequency Drives* (Bose, B. K. ed). IEEE Press, Piscataway, NJ, USA, pp. 138-208.

Horn, A., Pittorino, L.A. and Enslin, J.H.R. (1996): "Evaluation of Active Filter Control Algorithms under Non-Sinusoidal and Unbalanced Conditions". *Seventh International Conference on Harmonics and Quality of Power*, Las Vegas, USA, pp. 217-224.

Horowitz, P. and Hill, W. (1989): *The Art of Electronics*, Cambridge University Press, Cambridge, UK.

IEC 555-2 (1982): *Disturbances in Supply Systems Caused by Household Appliances and Similar Equipment*, International Electrotechnical Commission.

IEC 1000-4-7 (1991): *Electromagnetic compatibility (EMC) — Part 4: Testing and measurement techniques — Section 7: General guide on harmonics and interharmonics measurements and instrumentation, for power supply systems and equipment connected thereto*, International Electrotechnical Commission.

IEEE Std. 519 (1992): *IEEE Recommended Practices and Requirements for Harmonic Control in Electric Power Systems*, Institute of Electrical and Electronics Engineers.

IEEE Working Group on Power System Harmonics (1983): "Power Systems Harmonics: An Overview", *IEEE Transactions on Power Apparatus and Systems*, Vol. PAS-102, No. 8, pp. 2455-2460.

Ingram, D.M.E. and Round, S.D. (1997): "A Novel Digital Hysteresis Current Controller for an Active Power Filter". *Proceedings of Second International Conference on Power Electronics and Drive Systems (PEDS '97)*, Singapore, Vol. 2, pp. 744-749.

Jou, H.L. (1995): "Performance comparison of the three-phase active-power-filter algorithms", *IEE Proceedings. Generation, Transmission and Distribution*, Vol. 142, No. 6, pp. 646-652.

Kazmierkowski, M.P. and Dzieniakowski, M.A. (1994): "Review of Current Regulation Techniques For Three-Phase PWM Inverters". *20th International Conference on Industrial Electronics, Control and Instrumentation (IECON '94)*, Bologna, Italy, pp. 567-575.

Laird, H.D. (1992): "A Resonant DC Link Shunt Active Filter for the Power System", M.E. Thesis, University of Canterbury, Christchurch, New Zealand.

Lee, H., Suh, B.-S. and Hyun, D.-S. (1996): "A Novel PWM Scheme for a Three-Level Voltage Source Inverter with GTO Thyristors", *IEEE Transactions on Industry Applications*, Vol. 32, No. 2, pp. 260-268.

Li, Z., Jin, H. and Joos, G. (1995): "Control of Active Filters Using Digital Signal Processors". *21st International Conference on Industrial Electronics, Control and Instrumentation (IECON '95)*, Orlando, USA, Vol. 1, pp. 651-655.

Lin, K.-S., Frantz, G.A. and Simar, R. (1987): "The TMS320 Family of Digital Signal Processors", *Proceedings of the IEEE*, Vol. 75, No. 9, pp. 1143-1159.

Lorenz, R.D., Lipo, T.A. and Novotny, D.W. (1997): "Motion Control with Induction Motors", in *Power Electronics and Variable Frequency Drives* (Bose, B. K. ed). IEEE Press, Piscataway, NJ, USA, pp. 209-276.

Malesani, L., Rossetto, L., Tomasin, P. and Zuccato, A. (1996): "Digital Adaptive Hysteresis Current Control With Clocked Commutations and Wide Operating Range", *IEEE Transactions on Industry Applications*, Vol. 32, No. 2, pp. 316-325.

Malesani, L., Tenti, P., Gaio, E. and Piovan, R. (1991): "Improved Current Control Technique of VSI PWM Inverters with Constant Modulation Frequency and Extended Voltage Range", *IEEE Transactions on Industry Applications*, Vol. 27, No. 2, pp. 365-369.

Malesani, L. and Tomasin, P. (1993): “PWM Current Control Techniques of Voltage Source Converters - A Survey”. *19th International Conference on Industrial Electronics, Control and Instrumentation (IECON '93)*, Hawaii, USA, Vol. 2, pp. 670-675.

Mathworks (1996a): *Signal Processing Toolbox for MATLAB*, The Mathworks Inc.

Mathworks (1996b): *Simulink: Dynamic System Simulation for MATLAB*, The Mathworks Inc.

Maxim (1995): *Low-Power, 8-Channel, Serial 12-Bit ADCs*, Maxim Integrated Products, Sunnyvale, CA, USA.

Mazzocco, D. (1992): *Real Inverse FFT*, Texas Instruments, Houston, USA.

Morán, L.A., Fernández, L., Dixon, J.W. and Wallace, R.R. (1997): “A Simple and Low-Cost Control Strategy for Active Power Filters Connected in Cascade”, *IEEE Transactions on Industrial Electronics*, Vol. 44, No. 5, pp. 621-629.

Morán, L.A., Godoy, P., Wallace, R.R. and Dixon, J.W. (1993): “A New Current Control Strategy for Active Power Filters Using Three PWM Voltage Source Inverters”. *IEEE 24th Power Electronics Specialists Conference (PESC '93)*, Seattle, USA, pp. 3-9.

Nabae, A. and Tanaka, T. (1996): “A Universal Theory of Instantaneous Active-Reactive Current and Power Including Zero-Sequence Component”. *Seventh International Conference on Harmonics and Quality of Power*, pp. 90-95.

Pahmer, C., Capolino, G.A. and Henao, H. (1994): “Computer-Aided Design for Control of Shunt Active Filter”. *20th International Conference on Industrial Electronics, Control and Instrumentation (IECON '94)*, Bologna, Italy, pp. 669-674.

Poularikas, A.D., Bomar, B.W., Smith, L.M. and Cadzow, J.A. (1993): “Digital Signal Processing”, in *The Electrical Engineering Handbook Series* (Dorf, R. C. ed). CRC Press Inc, Boca Raton, Florida, USA, pp. 229-278.

Quinn, C.A., Mohan, N. and Mehta, H. (1993): “A Four-Wire, Current-Controlled Converter Provides Harmonic Neutralization in Three-Phase, Four-Wire Systems”. *Eight Annual Applied Power Electronics Conference and Exposition Proceedings*, San Diego, USA, pp. 841-846.

Rahman, K.M., Rezwan Khan, M., Choudhury, M.A. and Rahman, M.A. (1997): "Variable-Band Hysteresis Current Controllers for PWM Voltage-Source Inverters", *IEEE Transactions on Power Electronics*, Vol. 12, No. 6, pp. 964-970.

Redl, R., Tenti, P. and van Wyk, J.D. (1997): "Power electronics' polluting effect", *IEEE Spectrum*, Vol. 34, No. 5, pp. 32-39.

Retif, J.M., Allard, B., Jorda, X. and Perez, A. (1993): "Use of ASICs in PWM Techniques for Power Converters". *19th International Conference on Industrial Electronics, Control and Instrumentation (IECON '93)*, Maui, Hawaii, USA, Vol. 2, pp. 683-688.

Rice, D.E. (1986): "Adjustable Speed Drives and Power Rectifier Harmonics - Their Effect on Power Systems Components", *IEEE Transactions on Industry Applications*, Vol. IA-22, No. 1, pp. 161-177.

Round, S.D. (1992): "Real Time Optimisation of an Active Filter's Performance and Applications to the Power System", Ph.D Thesis, University of Canterbury, Christchurch, New Zealand.

Round, S.D. and Duke, R.M. (1994): "Real-Time Optimization of an Active Filter's Performance", *IEEE Transactions on Industrial Electronics*, Vol. 41, No. 3, pp. 278-284.

Round, S.D. and Ingram, D.M.E. (1997): "An Evaluation of Techniques for Determining Active Filter Compensating Currents in Unbalanced Systems". *7th European Conference on Power Electronics and Applications (EPE '97)*, Trondheim, Norway, Vol. 4, pp. 4.767-4.772.

Sasaki, H. and Machida, T. (1971): "A New Method to Eliminate AC Harmonic Currents by Magnetic Flux Compensation - Considerations on Basic Design", *IEEE Transactions on Power Apparatus and Systems*, Vol. PAS-90, No. 5, pp. 2009-2019.

Siemens (1996): *Power Semiconductor IGBT Modules*, Siemens Semiconductor Group.

Simar, R. and Davis, A. (1988): "The Application of High-Level Languages to Single-Chip Digital Signal Processors". *1988 International Conference on Acoustics, Speech and Signal Processing*, New York City, Vol. 3, pp. 1678-1681.

Stanislawski, J., Heydt, G., Enjeti, P., Steffek, L., *et al.* (1997): "Main Disturbances", in *The Industrial Electronics Handbook* (Irwin, J. D. ed). CRC Press LLC, Boca Raton, Florida, USA, pp. 349-381.

Tessarolo, A. (1991): *Fast FFT Algorithm for TMS320C30*, Texas Instruments, Sydney, Australia.

Texas Instruments (1994): *TMS320C3x User's Guide*, Texas Instruments Inc., Austin, Texas, USA.

Tzou, Y.-Y. and Hsu, H.-J. (1997): "FPGA Realization of Space-Vector PWM Control IC for Three-Phase PWM Inverters", *IEEE Transactions on Power Electronics*, Vol. 12, No. 6, pp. 953-963.

van Wyk, J.D. (1993): "Power Quality, Power Electronics and Control". *Fifth European Conference on Power Electronics and Applications*, Brighton, UK, IEE, Vol. 1, pp. 17-32.

Xilinx (1996): *The Programmable Logic Data Book*, Xilinx, San Jose, California, USA.

Appendix A
Circuit Schematics

A.1 Load Current ADC Board

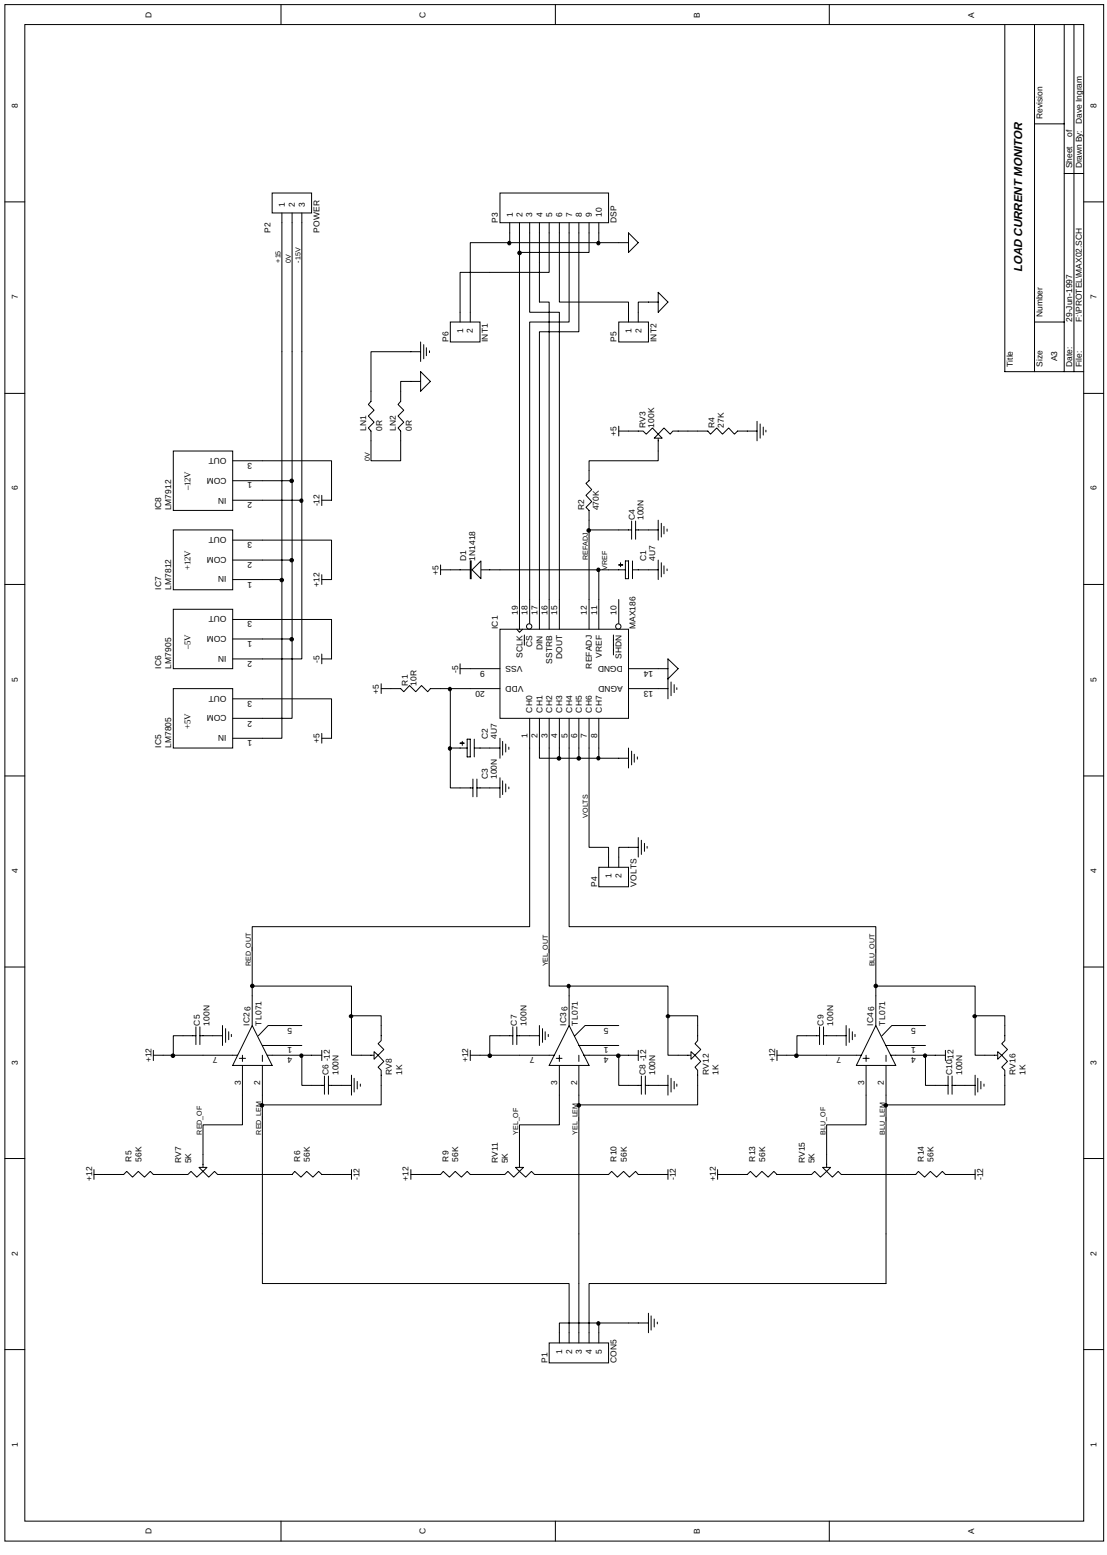

A.2 Injection Current ADC Board

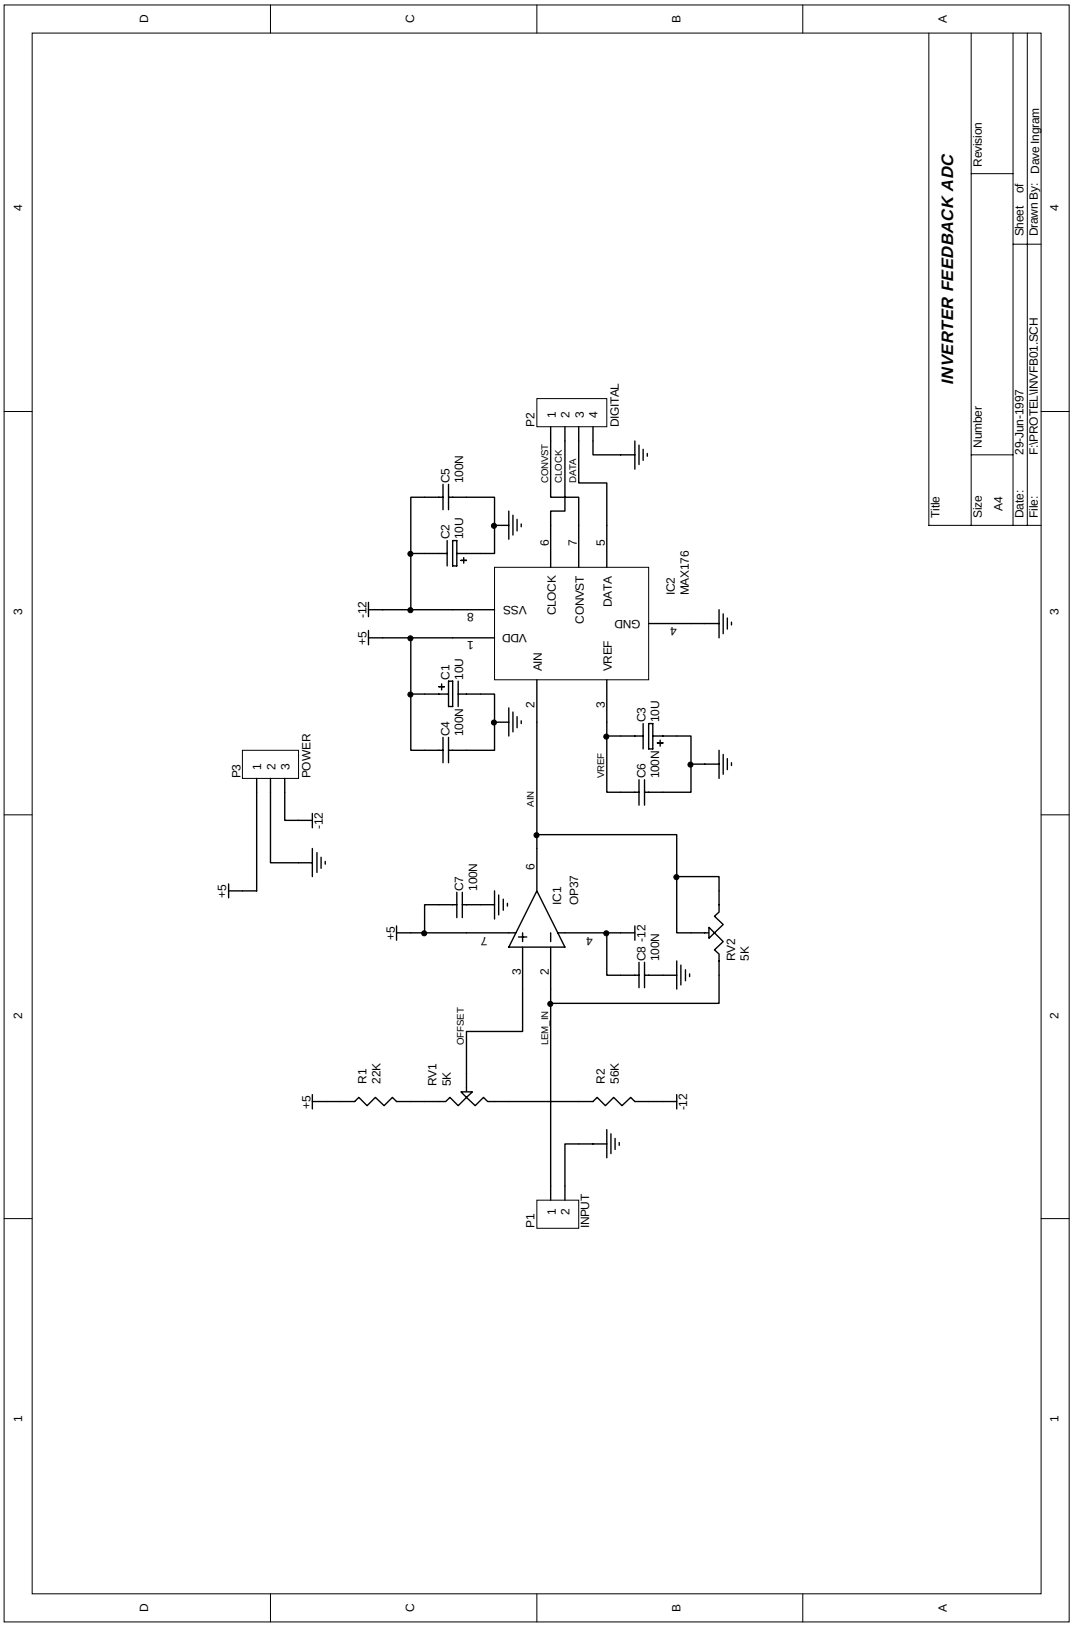

A.3 Sample Rate Multiplier

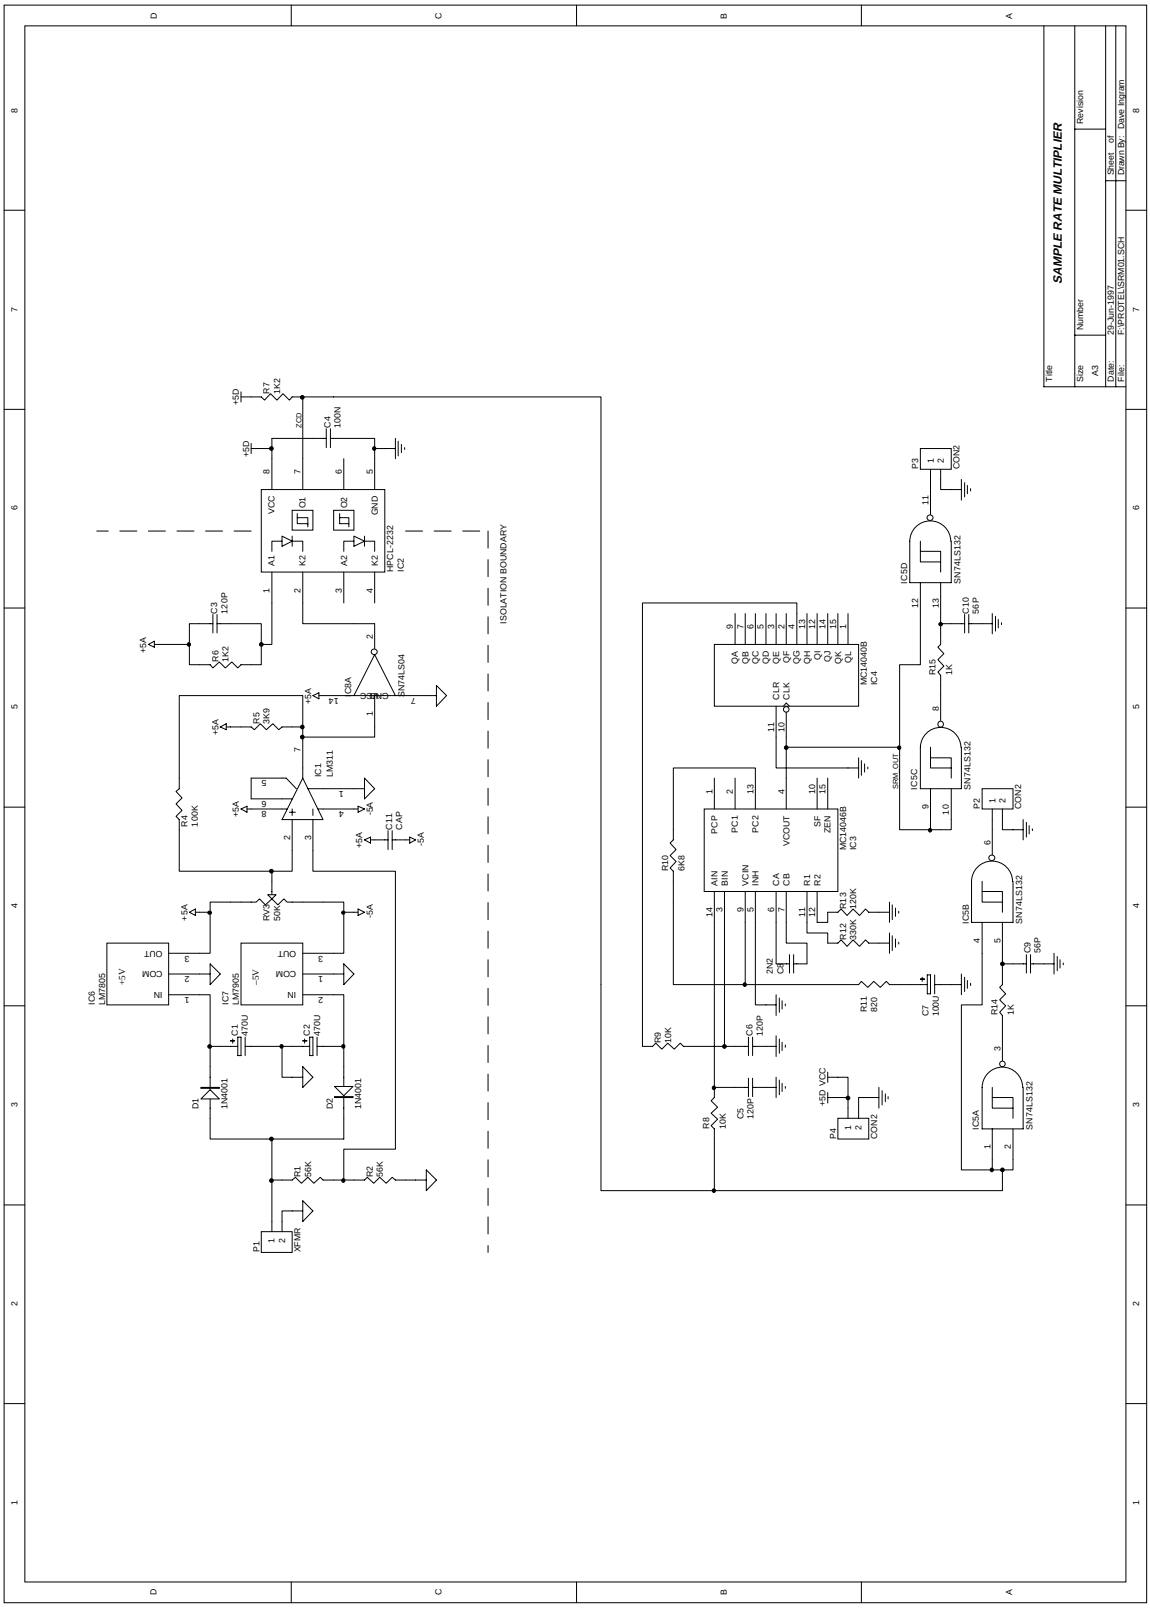

A.4 Three Phase Inverter

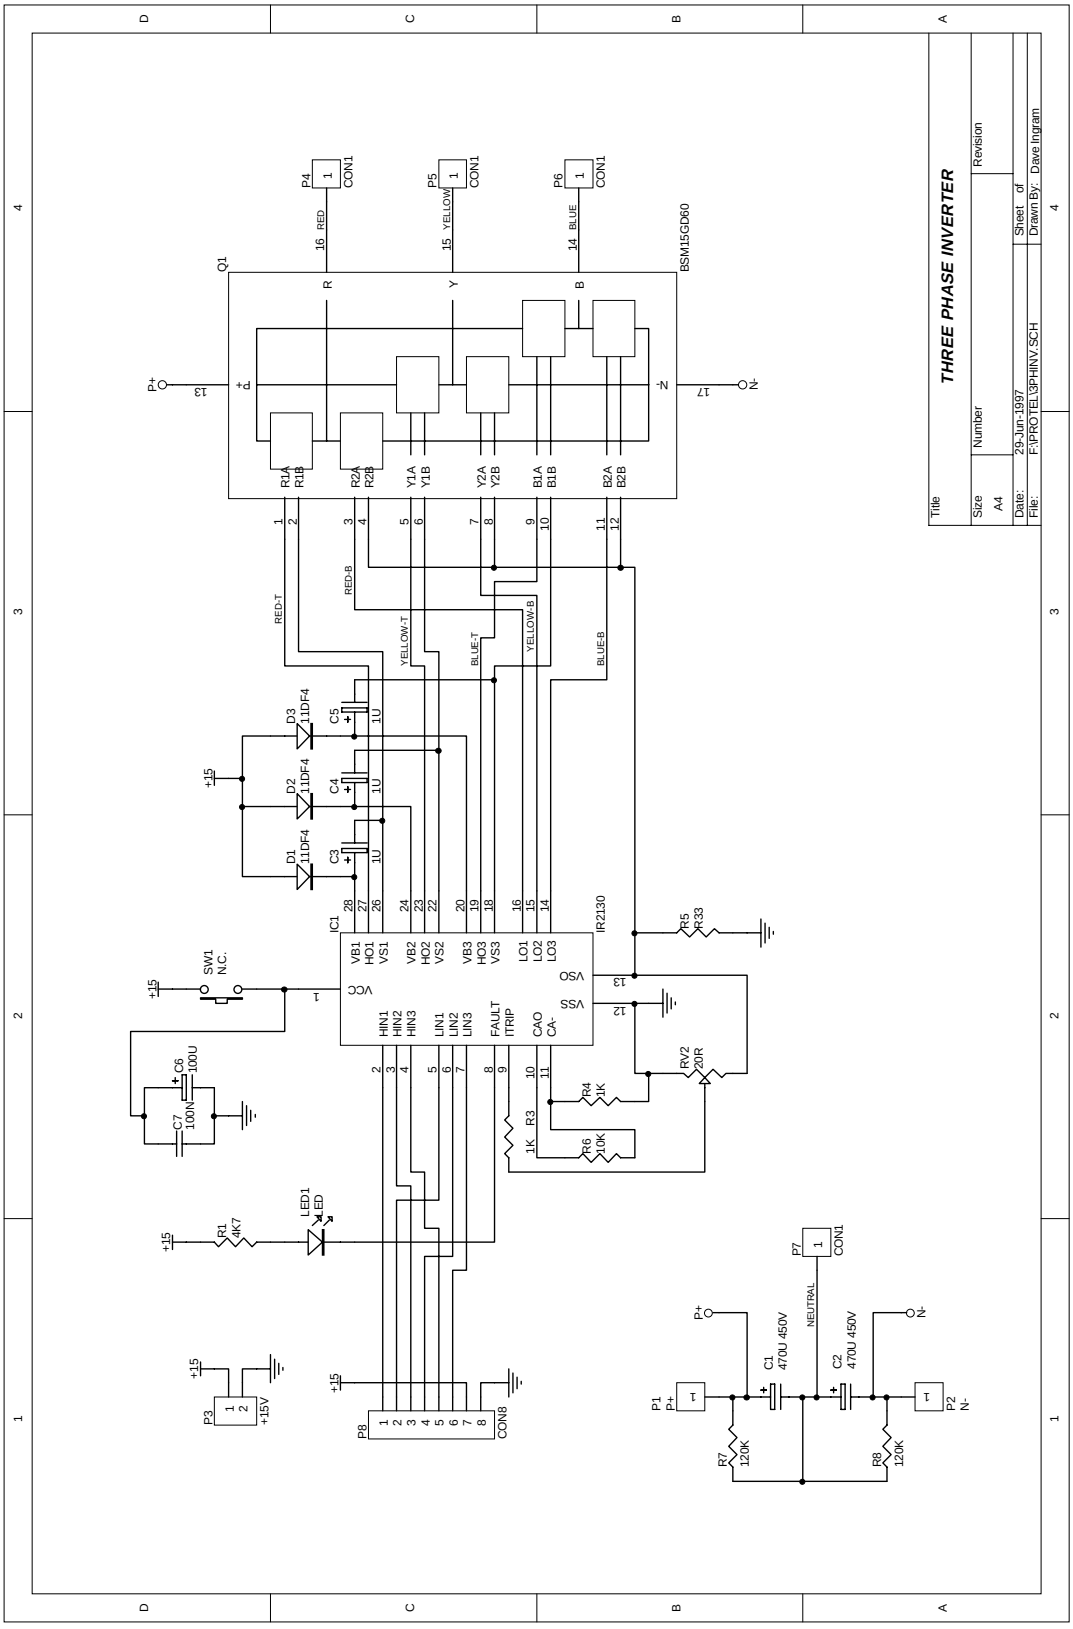

Appendix B
Xilinx Schematics

B.1 Digital Hysteresis Current Controller Top Level

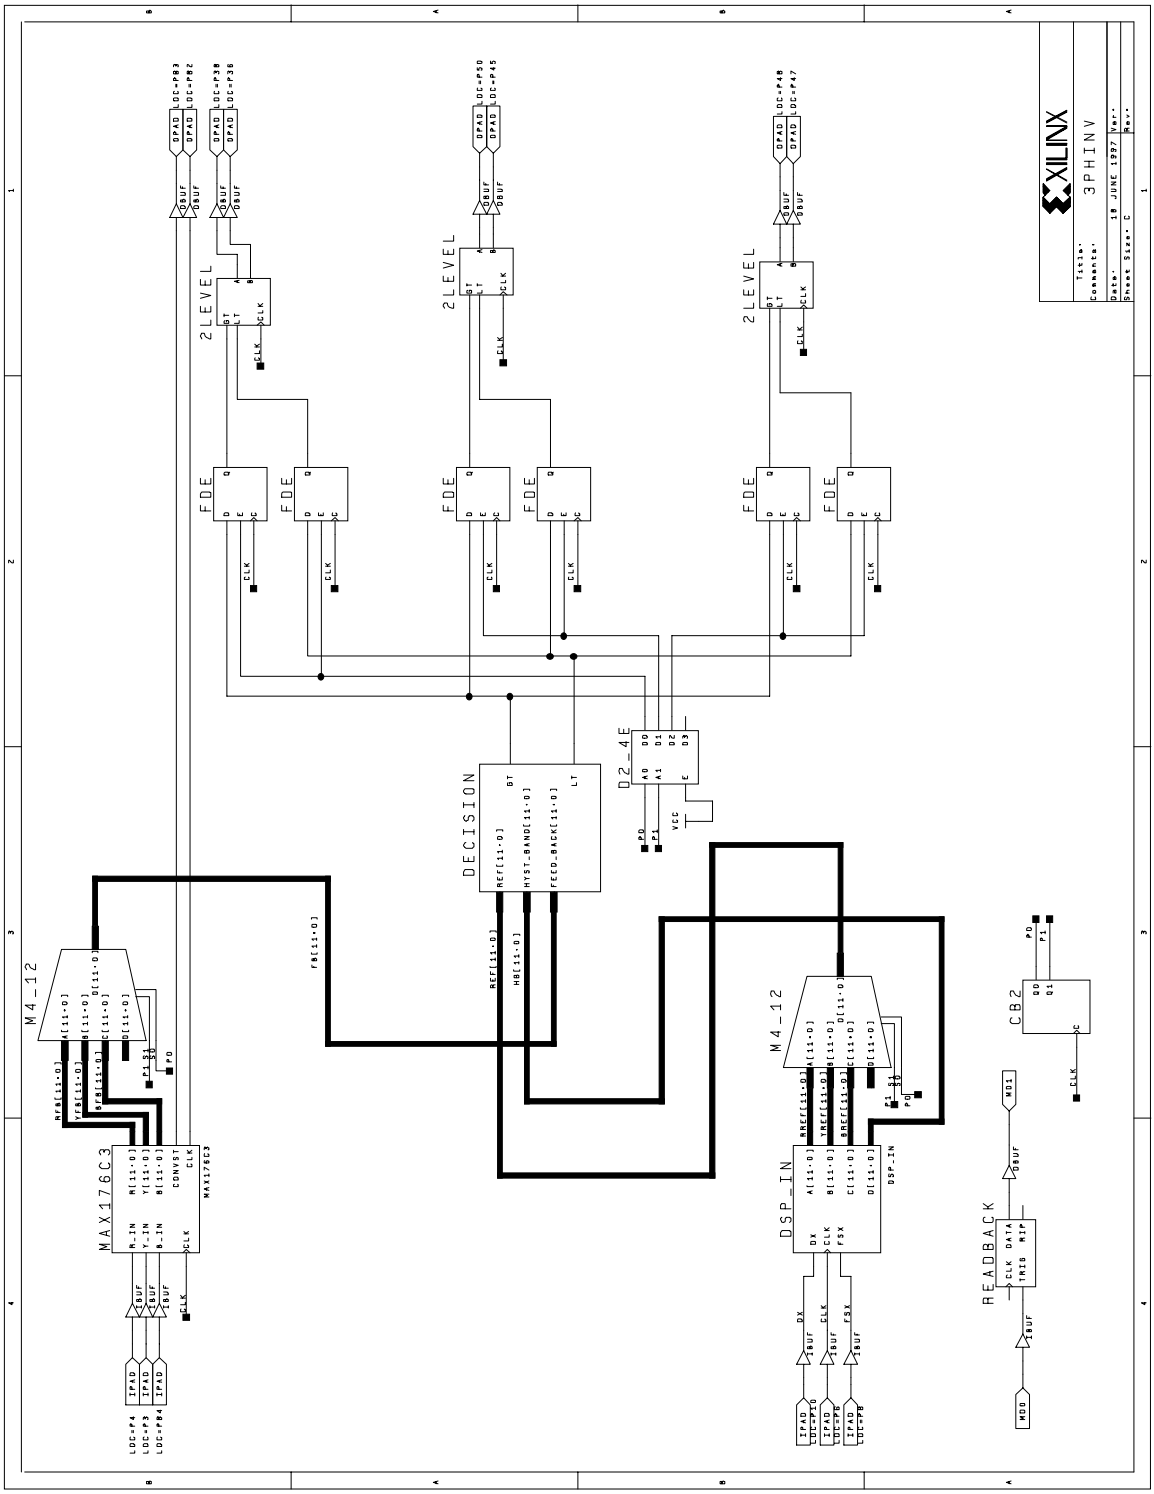

B.2 MAX176 Interface

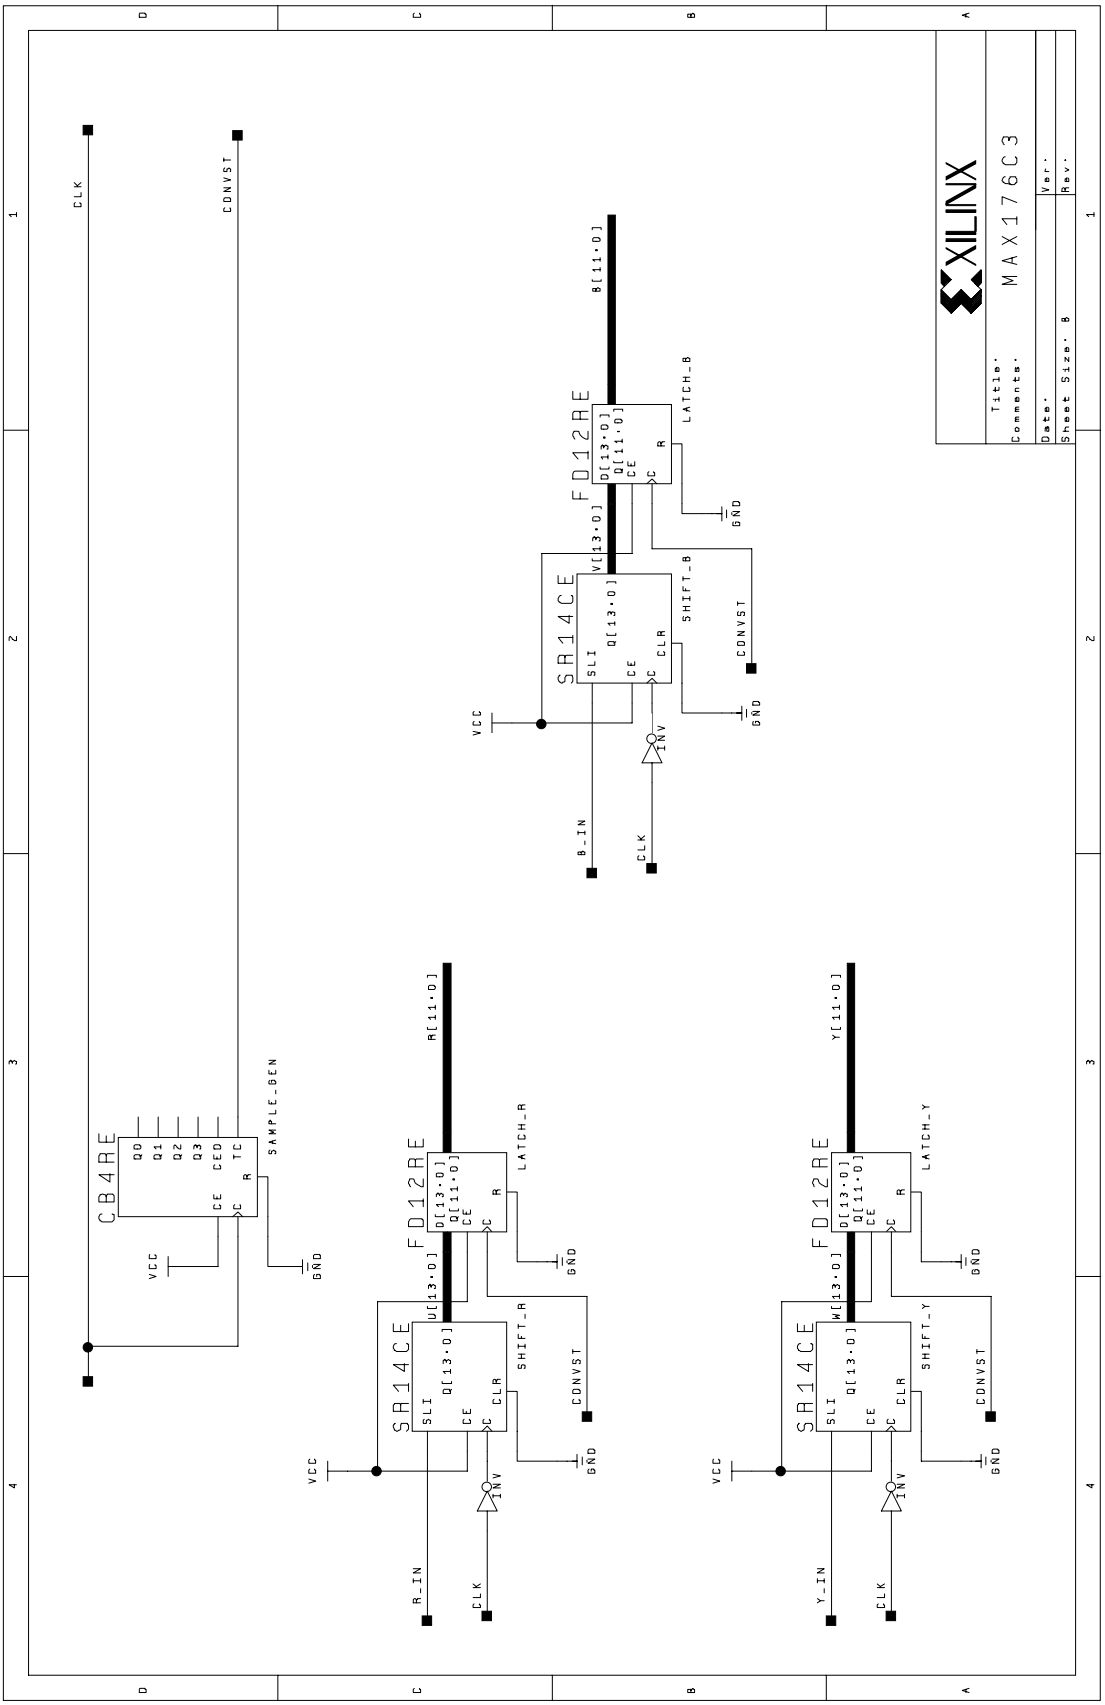

B.3 DSP Interface

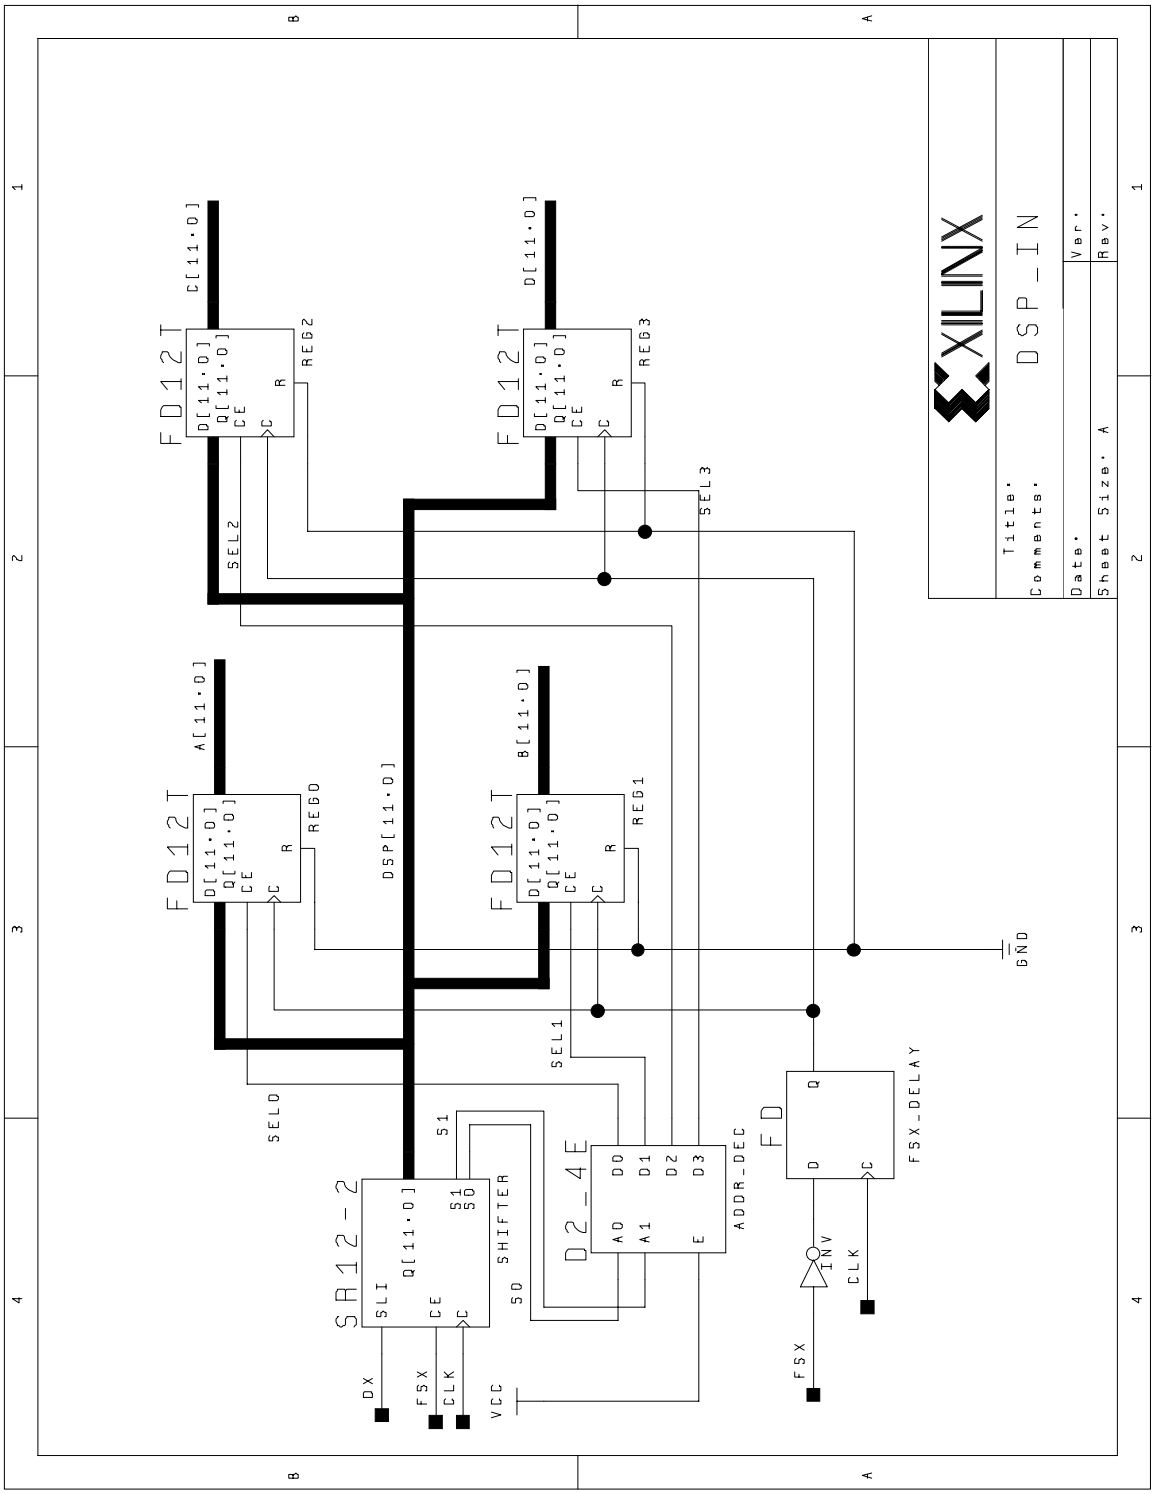

B.4 Decision Making Logic

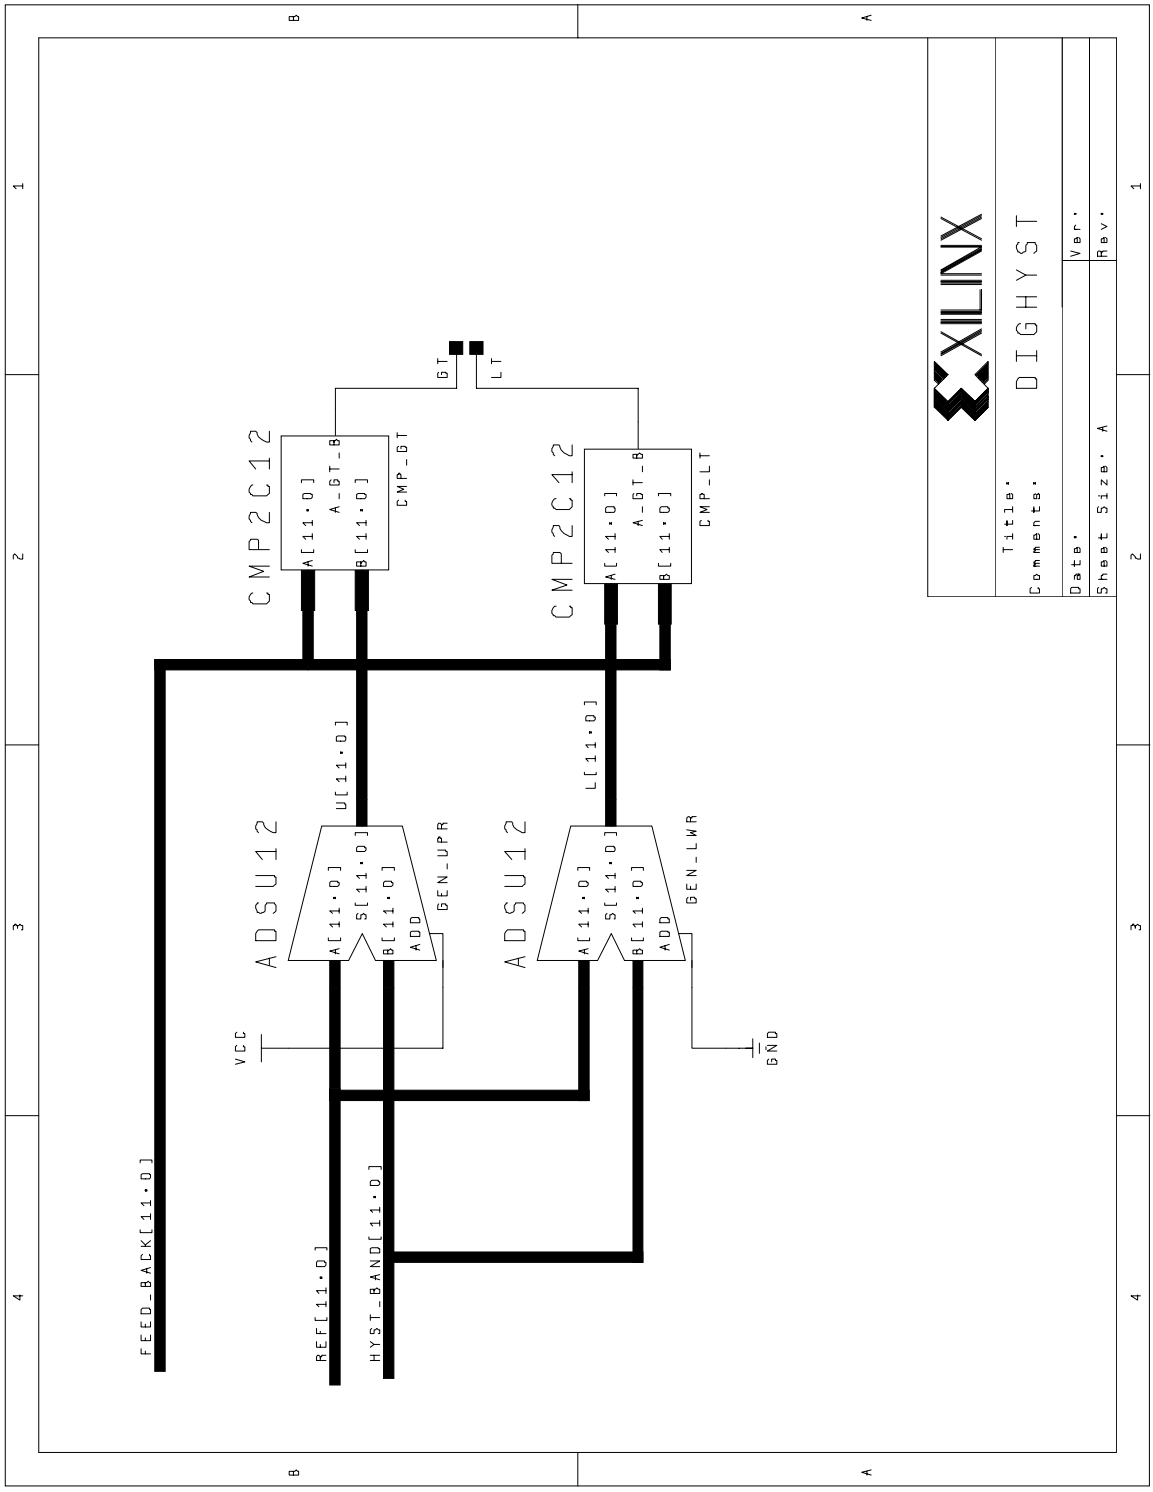

B.5 Deadtime Protection

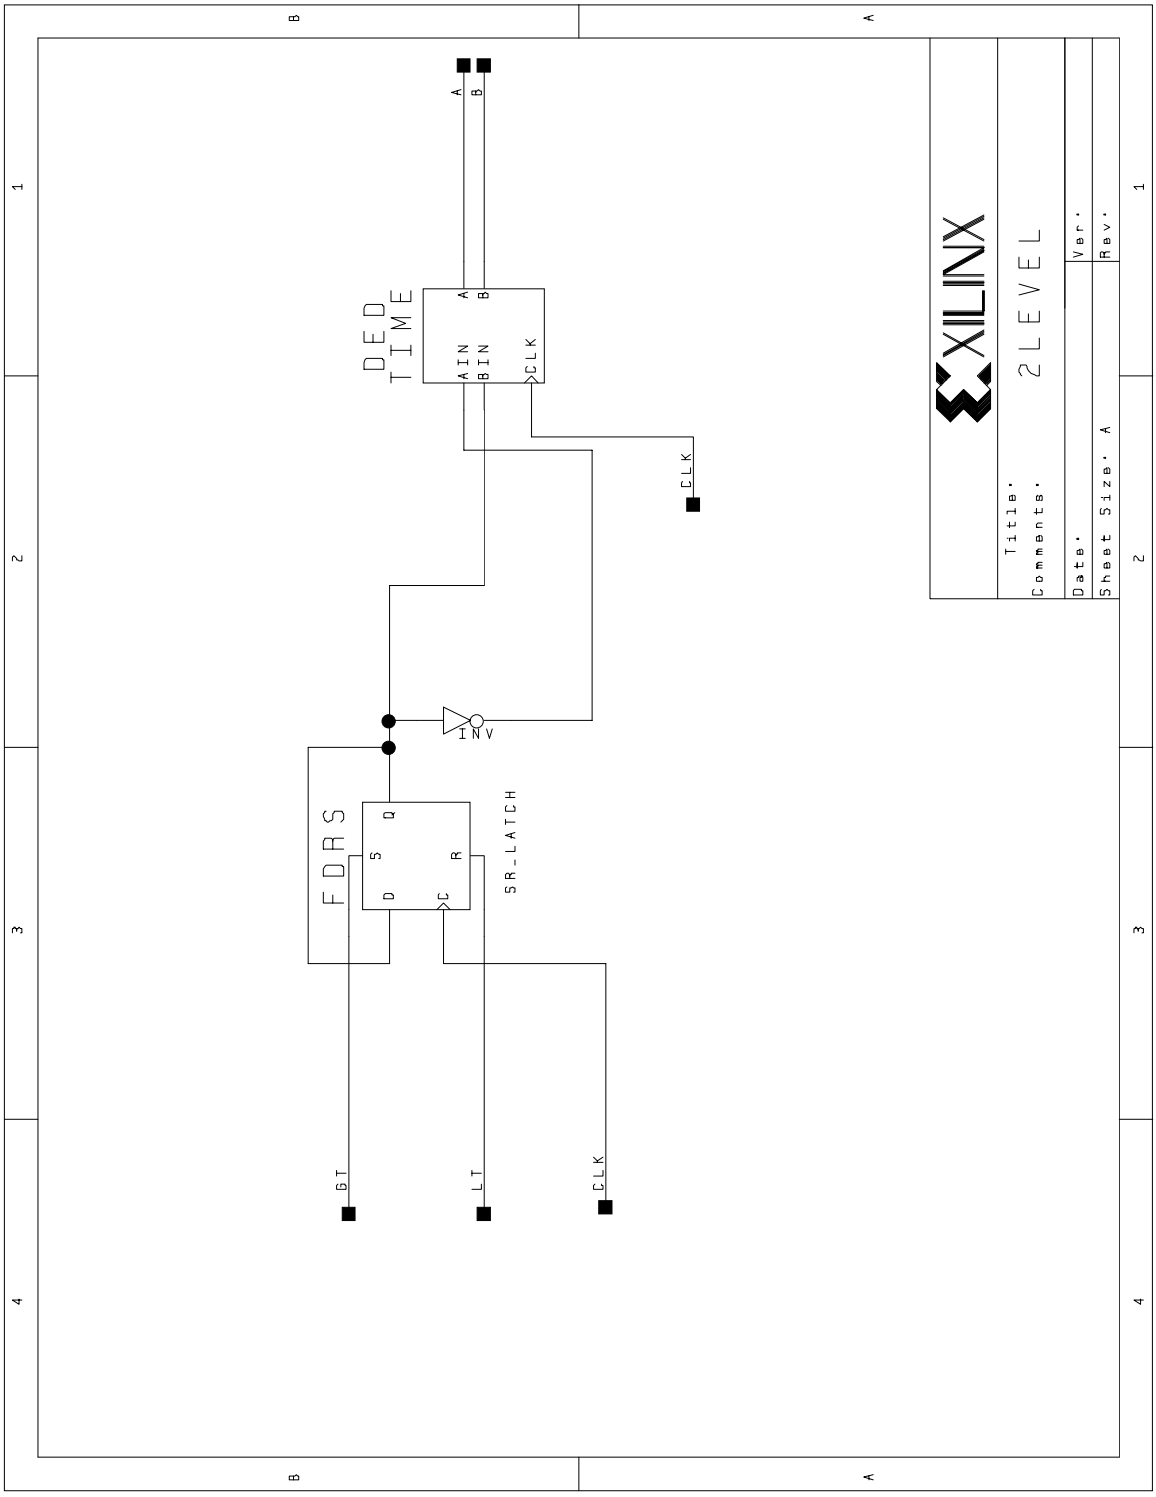

B.6 Dead Time State Machine

```
MODULE wait13
TITLE  '13 CLK Wait'

"Clock
       clk    PIN;

"Input
       hold   PIN;

"Output
       o      PIN;

"State Diagram Declaration and Assignment
       waitreg STATE_REGISTER ISTYPE 'REG_D';
       s0,s1,s2,s3,s4,s5,s6,s7,s8,s9,s10,s11,s12,s13 STATE;
       xilinx property 'INITIALSTATE s0';

Equations
       waitreg.clk = clk;
       o = s13;

State_Diagram waitreg
       STATE s0: if hold then s0 else s1;
       STATE s1: if hold then s0 else s2;
       STATE s2: if hold then s0 else s3;
       STATE s3: if hold then s0 else s4;
       STATE s4: if hold then s0 else s5;
       STATE s5: if hold then s0 else s6;
       STATE s6: if hold then s0 else s7;
       STATE s7: if hold then s0 else s8;
       STATE s8: if hold then s0 else s9;
       STATE s9: if hold then s0 else s10;
       STATE s10: if hold then s0 else s11;
       STATE s11: if hold then s0 else s12;
       STATE s12: if hold then s0 else s13;
       STATE s13: if hold then s0 else s13;

end
```